


THE
**INHABITED
BODY**

The Science of Our Microbial Selves



Dr Horst Herb

BACTERIA · ARCHAEA · EUKARYA

The Inhabited Body — Volume 1: The Other Us

A Journey Through the Human Microbiome

Horst Herb

2026-04-14

Contents

The Inhabited Body — Volume 1: The Other Us	6
Chapter 1: You Are Never Alone	7
1.1 The “10:1” Myth and the Real Numbers	7
1.2 What Do We Mean by “Microbiome”?	9
1.3 A Brief History of Germ Theory and How It Shaped Our Adversarial View of Microbes	9
1.4 The Paradigm Shift: From Enemies to Ecosystem	11
1.5 A Tourist’s Map of the Human Microbiome	12
You Are Never Alone — And You Wouldn’t Want to Be	14
Chapter 2: Built Together — How Microbes Shaped Human Evolution	16
2.1 The First Great Merger: Mitochondria and the Birth of Complexity	16
2.2 The Viral Archive: Eight Per Cent of You Is Virus	17
2.3 The Gift of the Placenta: Syncytins and the Viral Origins of Pregnancy	18
2.4 How a Virus Wired the Nervous System: RetroMyelin and the Insulation of Thought	20
2.5 How a Virus Gave Us Memory: The ARC Protein	22
2.6 How a Transposon Built the Immune System: RAG and V(D)J Recombination	23
2.7 Not Junk: Transposable Elements as an Evolutionary Toolkit	24
2.8 Bacterial Genes in the Human Genome	25
2.9 Evolution as Collaboration	25
Chapter 3: Your First Inheritance — Colonisation Across a Lifetime	28
3.1 Before Birth: The Sterile Womb Debate	28
3.2 The Birth Canal: Your First Microbial Baptism	29
3.3 Breastmilk: The Microbial Care Package	30
3.4 The First Thousand Days: Assembling a Stable Ecosystem	32
3.5 The Adult Microbiome: Stable but Not Static	33
3.6 Ageing and the Microbiome: The Final Transition	34
3.7 Your Microbiome Is Not Just Yours	35
Chapter 4: Tools of Discovery — How We Study What We Can’t See	38
4.1 From Petri Dishes to DNA: The Culture Problem	38
4.2 The Molecular Rosetta Stone: 16S Ribosomal RNA	39
4.3 The First Census: 16S Sequencing in Practice	40

4.4 Beyond the Barcode: Shotgun Metagenomics	41
4.5 What Are They Actually Doing? Metatranscriptomics, Metabolomics, and Proteomics	43
4.6 ITS Sequencing and the Mycobiome Blind Spot	44
4.7 Gnotobiotic Animals: What Happens When You Remove All the Microbes?	44
4.8 Computational Challenges: Bioinformatics, Big Data, and the Role of AI	46
4.9 The Sequencing Revolution: From Months to Minutes	47
4.10 The Return of Culture: Culturomics	48
4.11 Limitations and Pitfalls: How to Read Microbiome Science Critically	48
4.12 What the Tools Have Shown Us — and What Remains to Be Seen	51
Chapter 5: The Gut — Our Largest Microbial Organ	53
5.1 Anatomy of a Habitat: Stomach, Small Intestine, Colon — Different Worlds	53
5.2 The Dominant Phyla and What Drives Community Structure	55
5.3 The Mucus Layer: A Negotiated Border	56
5.4 Short-Chain Fatty Acids and Microbial Metabolism — The Currency of Cooperation	58
5.5 The Enteric Nervous System: A “Second Brain” Wired Into the Microbiome	60
5.6 Enterotypes — Do They Exist, and Do They Matter?	61
5.7 Colonisation Resistance — Why a Full House Is Hard to Invade	62
Chapter 6: The Skin — Life on the Surface	66
6.1 Deserts, Rainforests, and Caves: Ecological Niches Across the Skin	66
6.2 Residents vs Transients — Who Belongs and Who’s Passing Through	68
6.3 The Skin Barrier and Microbial Defence	69
6.4 <i>Malassezia</i> and the Skin Mycobiome — The Dominant Fungal Genus Hiding in Plain Sight	71
6.5 Acne, Eczema, and Wound Healing — When the Ecosystem Breaks Down	72
6.6 The Hygiene Hypothesis Revisited	75
Chapter 7: The Respiratory Tract — Not So Sterile After All	78
7.1 Overturning the Sterile Lung Dogma	78
7.2 The Nasopharynx as a Microbial Gateway	81
7.3 The Lung Microbiome in Health and Disease	82
7.4 The Lung Mycobiome: <i>Aspergillus</i> in Cystic Fibrosis and Post-COVID Fungal Super-infections	86
7.5 Lessons from COVID-19: Viral Infection and Microbial Disruption	88
Chapter 8: The Urogenital Microbiome	93
8.1 The Vaginal Microbiome — <i>Lactobacillus</i> Dominance and Community State Types	93
8.2 The Urinary Microbiome — Another “Sterile” Myth Debunked	97
8.3 Implications for Fertility, Pregnancy, and Preterm Birth	99
8.4 Sex Differences in Microbiome Composition and Disease Risk	102
Chapter 9: The Oral Microbiome — More Than Just Teeth	106
9.1 Biofilms and Dental Plaque: The Original Microbiome Research	106
9.2 Periodontal Disease and Systemic Health — The Cardiovascular Connection	110
9.3 The Oral-Gut Axis: Swallowing a Trillion Microbes a Day	113
9.4 The Oral Microbiome as a Diagnostic Window	116
Chapter 10: The Mycobiome — The Forgotten Kingdom	122

10.1 Why Fungi Were Overlooked: Technical Challenges in Sequencing and Culture . . .	122
10.2 Who's There? The Major Players	124
10.3 Scale and Proportion: A Minority by Numbers, Not by Influence	126
10.4 Fungal-Bacterial Interactions: Competition, Cooperation, and Cross-Kingdom Signalling	127
10.5 The Mycobiome and the Immune System: Dectin-1, CARD9, and Anti-Fungal Immunity as a Tonic Signal	129
10.6 <i>Candida</i> Expansion During Antibiotic Use — The Opportunist's Playbook	130
10.7 Gut Fungi and IBD: The Interkingdom Inflammation Loop	132
10.8 The Skin Mycobiome: <i>Malassezia</i> in Seborrheic Dermatitis, Dandruff, and Beyond .	134
10.9 The Infant Mycobiome: Colonisation, Maturation, and Early Immune Priming . . .	135
Chapter 11: The Brain Microbiome — The Final Frontier	139
11.1 The Dogma of the Sterile Brain and How It Fell	140
11.2 Evidence for Microbial Presence in the CNS — What We Know So Far	142
11.3 Methodological Controversies: Contamination Artefact or Genuine Residents? . . .	145
11.4 Potential Implications for Neurodegeneration	147
11.5 What Would a Brain Microbiome Actually <i>Do</i> ?	149
Chapter 12: The Immune System — A Lifelong Negotiation	154
12.1 Training the Immune System: How Microbes Teach Self vs Non-Self	155
12.2 Toll-Like Receptors, Pattern Recognition, and the Molecular Handshake	158
12.3 Regulatory T Cells and Oral Tolerance — Learning Not to Fight	161
12.4 Segmented Filamentous Bacteria and Th17 — A Case Study in Immune Shaping . .	164
12.5 When Negotiation Fails: Autoimmunity, Allergy, and Chronic Inflammation	166
12.6 The Microbiome and Cancer Immunotherapy — Why Some Patients Respond and Others Don't	169
A Negotiation Without End	171
Chapter 13: The Gut-Brain Axis — Microbes and the Mind	176
13.1 The Vagus Nerve as an Information Superhighway	176
13.2 Microbial Production of Neurotransmitters: Serotonin, GABA, Dopamine	179
13.3 The HPA Axis, Stress, and Microbial Modulation of Cortisol	182
13.4 Animal Studies: Anxious Mice, Bold Mice, and Faecal Transplants	184
13.5 Depression, Anxiety, and the Emerging Concept of "Psychobiotics"	187
13.6 Autism Spectrum Disorder — Signal in the Noise?	190
The Conversation Continues	193
Chapter 14: Behavioural Modulation — Are We the Puppet or the Puppeteer?	196
14.1 Can Microbes Influence Food Cravings?	196
14.2 Social Behaviour and Microbial Signalling in Animal Models	200
14.3 The Evolutionary Logic: Microbial Manipulation for Transmission Advantage . . .	203
14.4 Free Will, Determinism, and the Uncomfortable Philosophical Implications	206
14.5 How Seriously Should We Take the "Microbes Control Behaviour" Hypothesis? . . .	207
Chapter 15: Metabolism, Hormones, and the Endocrine Microbiome	211
15.1 Bile Acid Metabolism and the Gut-Liver Axis	212
15.2 The Estrobolome — Microbial Regulation of Oestrogen	215
15.3 Thyroid Function and the Microbiome	218

15.4 Circadian Rhythms — Do Our Microbes Have Body Clocks?	221
The Hormonal Ecosystem	224
Chapter 16: Bacteriophages — The Viruses Within the Microbiome	227
16.1 What Are Phages and Why Were They Forgotten?	228
16.2 The Human “Phageome” — Scale, Diversity, and Dynamics	231
16.3 Phages as Regulators of Bacterial Populations	234
16.4 Lysogeny: When Viruses Become Part of the Bacterial Genome	236
16.5 Phage-Mediated Horizontal Gene Transfer — Shuffling the Deck	238
The Missing Players	240
Chapter 17: Phages, Immunity, and the Tripartite Relationship	243
17.1 Phage Transcytosis Across Epithelial Barriers	243
17.2 Direct Interactions Between Phages and the Mammalian Immune System	245
17.3 The BAM Model — A Non-Host Immunity Layer	248
17.4 Phages as Anti-Inflammatory Agents? Emerging Evidence	251
The Complete Ecosystem	252
The Inhabited Body — Bibliography by Chapter	257
Ch. 1 — You Are Never Alone	257
Ch. 2 — Built Together: How Microbes Shaped Human Evolution	257
Ch. 3 — Your First Inheritance: Colonisation Across a Lifetime	258
Ch. 4 — Tools of Discovery	258
Ch. 5 — The Gut	258
Ch. 6 — The Skin	260
Ch. 7 — The Respiratory Tract	260
Ch. 8 — The Urogenital Microbiome	262
Ch. 9 — The Oral Microbiome	263
Ch. 10 — The Mycobiome	263
Ch. 11 — The Brain Microbiome	264
Ch. 13 — The Gut-Brain Axis	264
Ch. 14 — Behavioural Modulation	265
Ch. 15 — Metabolism, Hormones, and the Endocrine Microbiome	265
Ch. 16–17 — Bacteriophages	265
Ch. 18 — Phage Therapy	266
Ch. 19 — Antibiotics: Collateral Damage	266
Ch. 20 — Diet: Feeding the Ecosystem	267
Ch. 21 — Environment, Lifestyle, and the Modern World	269
Ch. 22 — Drugs Beyond Antibiotics	270
Ch. 23 — Probiotics: Promise, Hype, and Hard Truths	273
Ch. 24 — Faecal Microbiota Transplantation	274
Ch. 26 — Inflammatory and Autoimmune Disease	274
Ch. 27 — Metabolic Disease and Obesity	274
Ch. 28 — Cancer	275
Ch. 29 — Neurological and Psychiatric Disease	275
Ch. 31 — Ethical, Social, and Philosophical Questions	276
Notable Observations	276
Ch. 13 — The Gut-Brain Axis: Microbes and the Mind	276

Ch. 17 — Phages, Immunity, and the Tripartite Relationship 277

The Inhabited Body — Volume 1: The Other Us

A Journey Through the Human Microbiome

Horst Herb

❑ **PROOF COPY — 2026-04-14 — NOT FOR DISTRIBUTION**

Chapter 1: You Are Never Alone

You are, at this very moment, teeming with life that is not your own.

As you read this sentence, several trillion microorganisms are going about their business inside and upon your body. Bacteria line the walls of your intestines in communities so dense they form a living carpet. Fungi nestle into the creases of your skin. Viruses - most of them not infecting *you* at all, but the bacteria that live alongside you - drift through your gut in numbers that beggar the imagination. Archaea, those ancient single-celled organisms that predate complex life by billions of years, quietly metabolise gases in your colon. And all of this is happening without your knowledge, without your permission, and - here is the part that matters - very much to your benefit.

You have never been alone. Not for a single moment of your life. From the instant you passed through the birth canal, or were lifted from your mother's abdomen, you began acquiring passengers. They colonised your skin, your mouth, your gut. They set up shop. And then, in ways that science is only now beginning to understand, they got to work — training your immune system, digesting your food, manufacturing vitamins your own cells cannot produce, and communicating with your brain through chemical signals that influence your mood, your appetite, and possibly even your behaviour.

This is not science fiction. This is not a metaphor. This is you.

The community of microorganisms that lives in and on every human being — collectively known as the *microbiome* — is one of the most important discoveries in modern biology. Not because these organisms are new; they have been with us for as long as “us” has existed. But because we have only recently developed the tools to see them, count them, and begin to understand what they are doing. And what they are doing turns out to be far more consequential than anyone suspected.

This book is about that discovery and what it means: for medicine, for our understanding of disease, for the way we eat, the drugs we take, and the chemicals we surround ourselves with. It is also about the ways we are getting it wrong: the well-intentioned interventions that backfire, the popular remedies that may cause more harm than good, and the sheer hubris of imagining that we can manipulate an ecosystem we barely understand.

But before we get to any of that, we need to answer a simpler question — one that turns out to be surprisingly hard to pin down.

How many of *them* are there?

1.1 The “10:1” Myth and the Real Numbers

For decades, a single statistic dominated every popular account of the human microbiome: the claim that microbial cells in the human body outnumber human cells by a ratio of ten to one. It was a marvellous number. It appeared in textbooks, in TED talks, in newspaper headlines. It was the kind of fact that made you feel pleasantly unsettled — the idea that you were, in some fundamental sense, more microbe than human.

The only problem was that it was wrong.

The figure's precise origins are murky, but early estimates from the 1970s (mainly based on a back-of-the-envelope calculation by the microbiologist Thomas Luckey) suggested that the human gut alone contained around 100 trillion (10^{14}) microbial cells. Since the standard estimate of cells in the human body at the time was around 10 trillion (10^{13}), a tidy ten-to-one ratio fell out of the arithmetic. The number was picked up by reviews and textbooks, repeated so often and by so many credible sources that no one thought to check the underlying assumptions for over forty years. [REF:luckey1972] [REF:berg1996] [REF:bianconi2013]

In 2016, a team led by Ron Sender, Shai Fuchs, and Ron Milo at the Weizmann Institute of Science in Israel decided to do exactly that. Their paper, published in *PLOS Biology*, was a meticulous re-examination of every estimate in the chain. They revisited the volume of the colon, the density of bacteria in faecal samples, the actual number of human cells (which, it turns out, had been substantially underestimated), and the contribution of different body sites. Their conclusion was striking: the ratio of bacterial to human cells in a “reference man” — a 70-kilogram, 170-centimetre, 20-to-30-year-old male — is approximately 1.3 to one. [REF:sender2016]

Not ten to one. Roughly *one to one*.

Moreover, the total number of bacteria was closer to 39 trillion, not 100 trillion — and the number of human cells was closer to 30 trillion, not 10 trillion. [REF:sender2016] The margin of error was large enough that, at any given moment, you might be slightly more human than microbe, or slightly more microbe than human. After a large bowel movement, you might temporarily tip the balance in favour of your own cells. It was, as Sender and colleagues drily noted, a “revised estimate.”

Does it matter? In one sense, no. Whether the ratio is ten to one or one to one, the fundamental point remains: you are home to an enormous community of non-human life. Trillions of microbes — bacteria, archaea, fungi, viruses, and the occasional protist — live mainly on body surfaces and in body cavities, with the gut harbouring the vast majority of them. [REF:sender2016] [REF:berg2020] Their total biomass is often estimated at a few hundred grams, roughly the weight of a large apple. That may not sound like much set against a 1,300-gram brain or a 300-gram heart, but it is comparable to the weight of a kidney — and unlike a kidney, the microbiome contains not one genome but millions of them. [REF:gilbert2018]

But the revised number does matter in another sense, because it changes the metaphor. At ten to one, you could think of yourself as a vehicle for microbes — a walking apartment building providing shelter for a population that vastly outnumbers its landlord. At one to one, the metaphor shifts. You are not outnumbered. You are in a *partnership*. And like all partnerships, the balance of power depends not on head counts but on what each partner brings to the table.

What the microbiome brings to the table, as we shall see, is a genetic repertoire that dwarfs our own. The human genome contains roughly 19,000 to 20,000 protein-coding genes. The collective genomes of the microorganisms that inhabit a single human body — sometimes called the *meta-genome* — encode hundreds of thousands of genes, and when researchers pool data across many individuals to build population-level catalogues, the totals run into the millions. [REF:qin2010] [REF:tierney2019] [REF:almeida2021] Even at the scale of a single person, that is more than an order of magnitude more genetic information than your own DNA carries; across the species, the gap is vastly larger still.

Think about what that means. The vast majority of the genetic information associated with your body is not human. It belongs to your microbial passengers. And much of that information encodes enzymes, metabolic pathways, and signalling molecules that your own cells cannot produce. Strip

away the microbiome — remove every last bacterium, fungus, and virus — and you would not merely be “cleaner.” You would be missing essential biological capabilities. You would be, in a very real sense, incomplete.

1.2 What Do We Mean by “Microbiome”?

Before we go further, we should untangle a piece of terminology that causes confusion even among scientists.

The word *microbiome* was popularised by the Nobel laureate Joshua Lederberg in 2001, who used it to describe “the ecological community of commensal, symbiotic, and pathogenic microorganisms that literally share our body space.” Lederberg’s definition was deliberately inclusive: he wanted a word that captured the entire microbial world associated with the human body, regardless of whether individual species were helpful, harmful, or somewhere in between.

In practice, however, two related terms have come into common use, and they are not quite interchangeable.

The *microbiota* refers to the actual organisms — the living bacteria, archaea, fungi, protists, and viruses that inhabit a particular environment. When we say “the gut microbiota,” we mean the community of microbes physically present in the gastrointestinal tract.

The *microbiome*, in its more precise usage, refers to the entire habitat: the organisms *and* their collective genetic material, their metabolic products, and the environmental conditions they create and respond to. The microbiome is the ecosystem; the microbiota is the cast of characters.

In everyday conversation and in most popular writing — including, for the most part, this book — the two terms are used interchangeably. When precision matters, I will make the distinction. But for now, think of the microbiome as the whole show: the organisms, their genes, their chemistry, and the ceaseless molecular conversation they conduct with each other and with us.

One more point of terminology. You will sometimes encounter the phrase *commensal organism*. In ecology, commensalism describes a relationship in which one organism benefits while the other is neither helped nor harmed. For a long time, the bacteria that lived on and in us were assumed to be commensals: freeloaders who took advantage of the warm, nutrient-rich environment of the human body without contributing anything in return. As we shall see, this assumption was spectacularly wrong. Most of the organisms in the human microbiome are not commensals. They are *mutualists* — partners in a relationship where both parties benefit. The word “commensal” persists in the literature out of habit, but it understates the case considerably.

1.3 A Brief History of Germ Theory and How It Shaped Our Adversarial View of Microbes

To understand why the microbiome was overlooked for so long, we need to understand the intellectual tradition it was hidden behind: the germ theory of disease.

The idea that invisible living organisms cause illness is one of the great triumphs of modern science. Before the mid-nineteenth century, disease was attributed to “miasmas” — noxious airs arising

from rotting organic matter — or to imbalances of the body's four humours. Cholera epidemics were thought to be caused by foul-smelling winds. Wound infections were regarded as an inevitable consequence of surgery. Childbed fever, which killed new mothers by the thousand, was a mystery that physicians attributed to emotional disturbance, atmospheric conditions, or simply bad luck.

Then came the microbe hunters.

Antonie van Leeuwenhoek, a Dutch draper with a talent for grinding lenses, had first observed bacteria in the 1670s — he called them *animalcules* and found them in everything from rainwater to his own dental scrapings. But it took another two centuries before anyone connected these tiny organisms to disease. Louis Pasteur, through a series of elegant experiments in the 1860s, demonstrated that fermentation and putrefaction were caused by living microorganisms, not by spontaneous generation. Robert Koch, working in Germany in the 1870s and 1880s, developed the rigorous framework — now known as Koch's postulates — for proving that a specific microbe causes a specific disease. Koch identified the causative agents of anthrax, tuberculosis, and cholera, and in doing so, he transformed medicine.

The impact was immediate and enormous. Joseph Lister introduced antiseptic surgery. Public health authorities began chlorinating water supplies. Vaccines were developed. And, in the twentieth century, antibiotics arrived — first sulfonamides in the 1930s, then Alexander Fleming's penicillin, mass-produced from the mid-1940s onward. Together, these advances saved hundreds of millions of lives.

But they also bequeathed us a particular way of thinking about microorganisms — one that was, at its core, adversarial. Germ theory taught us that microbes are the enemy. That the goal of medicine is to find them and kill them. That a sterile environment is a safe environment. And that the fewer microorganisms you have in and on your body, the healthier you are likely to be.

This was not a subtle bias. It was the dominant paradigm. For more than a century, the language of microbiology was the language of warfare: *infection, invasion, defence, attack, resistance*. Bacteria were *germs*. The immune system was an *army*. Antibiotics were *magic bullets*. And hygiene — the relentless campaign to eliminate microbial life from our homes, our hospitals, our food, and our bodies — was presented as an unalloyed good. More clean was always better. Sterile was the ideal.

It is worth pausing to acknowledge just how successful this paradigm was. Infectious disease was the leading cause of death in the developed world at the start of the twentieth century. By the end of it, thanks largely to the application of germ theory, it had fallen to a distant also-ran, overtaken by heart disease, cancer, and stroke. The eradication of smallpox, the near-eradication of polio, the transformation of bacterial pneumonia from a death sentence to an inconvenience — these are achievements of staggering importance, and they all rest on the foundational insight that specific microbes cause specific diseases.

The problem was not that germ theory was wrong. It was that it was *incomplete*. By focusing exclusively on the tiny minority of microorganisms that cause disease, it blinded us to the vast majority that do not. And by framing the relationship between humans and microbes as fundamentally antagonistic, it made it almost impossible to see the partnership that was hiding in plain sight.

Consider the numbers. There are estimated to be around one trillion species of microorganism on Earth. Of these, the number that cause disease in humans is vanishingly small — perhaps a few hundred species are regular human pathogens. That is a fraction of a fraction of a fraction of microbial diversity. Yet for over a century, these few hundred species received almost all of the

scientific attention, while the trillions of non-pathogenic species — including the thousands that live in and on us — were treated as background noise.

The metaphor that best captures this blind spot is one borrowed from ecology. Imagine studying the Amazon rainforest but only paying attention to jaguars. Jaguars are important; they shape the ecosystem in profound ways. But if you described the Amazon as “a place full of jaguars,” you would miss the trees, the insects, the birds, the fungi, the river systems, the microbial communities in the soil — in short, you would miss the forest. That is, in essence, what happened in microbiology. We studied the jaguars and missed the forest.

1.4 The Paradigm Shift: From Enemies to Ecosystem

The first hints that the adversarial model was inadequate came from observations so mundane that they were easily dismissed.

Clinicians had long noticed, for example, that patients treated with broad-spectrum antibiotics sometimes developed secondary infections — particularly *Clostridioides difficile* colitis, a severe and sometimes fatal inflammation of the colon. The explanation, in retrospect, was obvious: the antibiotics had killed not just the target pathogen but the patient’s normal gut bacteria as well, creating a vacancy that *C. difficile*, which is resistant to many antibiotics, was ideally positioned to exploit. But this observation did not fit neatly into the germ theory framework. It suggested that the body’s resident microbes were doing something useful — that they were, in fact, *protecting* the host by occupying ecological niches that would otherwise be available to pathogens.

Similarly, paediatricians had observed that children raised on farms had lower rates of asthma and allergies than children raised in cities. Children who attended daycare early — and were therefore exposed to a wider variety of microorganisms at a young age — also seemed to be partially protected. These observations, formalised in 1989 by the epidemiologist David Strachan as the “hygiene hypothesis,” suggested that early microbial exposure was somehow important for the normal development of the immune system, and that too-clean environments might paradoxically increase the risk of immune disorders.

But the real revolution came with technology. The invention of culture-independent methods for identifying microorganisms — techniques that did not require growing bacteria in the laboratory but instead identified them directly from their DNA — revealed a microbial world of astonishing richness and complexity. When researchers began sequencing the genes of microorganisms recovered from the human gut, they discovered not a handful of species but *hundreds*, in some cases over a thousand, coexisting in intricate, interdependent communities. Many of these species had never been grown in a laboratory. Some had never been described. They were the dark matter of human biology — invisible to traditional methods but present in enormous numbers, performing functions that turned out to be essential.

The Human Microbiome Project, launched by the U.S. National Institutes of Health in 2007 and completed in its first phase in 2012, was the landmark effort that brought all of this into focus. By sampling 242 healthy adults at 18 body sites and performing both 16S ribosomal RNA sequencing and shotgun metagenomic analysis — technologies we will explore in Chapter 3 — the project produced the first comprehensive atlas of the healthy human microbiome. The results were revelatory. Every individual carried a unique microbial signature. The composition of the microbiome

varied enormously between body sites — the community in your mouth bore almost no resemblance to the community in your gut, which in turn was radically different from the community on your skin. And the functional capabilities encoded in the microbial metagenome were vast, encompassing metabolic pathways that human cells simply do not possess.

The paradigm began to shift. Microorganisms were no longer just pathogens to be eliminated; they were ecological partners to be understood. The language of warfare started giving way to the language of ecology: *community*, *diversity*, *symbiosis*, *ecosystem*, *niche*. And with this new language came a new set of questions. Not “how do we kill the bad microbes?” but “how does the whole community function?” Not “how do we sterilise the environment?” but “what happens when we disrupt the ecosystem?”

This shift in thinking carries a subtle but important implication that will recur throughout this book, and it is worth flagging now. If the old paradigm sorted microbes into “good” and “bad,” the new one recognises that these labels are almost always misleading. A species of bacterium that is beneficial in one context — occupying a niche in your gut, producing short-chain fatty acids that nourish the cells of your colon wall — can be harmful in another context, if it escapes the gut and enters the bloodstream. *Escherichia coli* is a normal and important resident of the human intestine; it is also the most common cause of urinary tract infections. *Staphylococcus epidermidis* lives harmlessly on the skin of virtually every human being; it is also one of the leading causes of infections on implanted medical devices. Context matters. Location matters. The state of the host’s immune system matters. The composition of the surrounding community matters.

This principle — that the same organism can be friend or foe depending on circumstances — has profound implications for how we think about interventions aimed at the microbiome. It means that adding a “good” bacterium to a system where it does not belong, or where it competes with organisms that are already performing essential functions, can cause harm rather than benefit. We will return to this idea repeatedly in the chapters that follow, because it is one of the most important — and most frequently misunderstood — lessons of microbiome science.

1.5 A Tourist’s Map of the Human Microbiome

Imagine, for a moment, that you could shrink yourself to the scale of a bacterium and take a tour of the human body. What would you find?

The answer depends entirely on where you look, because the human microbiome is not one ecosystem but many — a patchwork of distinct habitats, each with its own physical conditions, its own chemical environment, and its own characteristic community of residents. A bacterium adapted to the oxygen-poor, bile-rich environment of the large intestine would no more survive on the dry, acidic surface of the forearm than a deep-sea fish would thrive in a desert. The diversity of microbial life across the human body reflects the diversity of the environments the body provides.

The gut is the main event. The gastrointestinal tract, and the large intestine in particular, harbours the vast majority of the body’s microbial biomass — somewhere between 70 and 80 per cent of all bacteria in the body reside here. The colon alone contains bacterial densities of up to 100 billion cells per gram of content, making it one of the most densely populated microbial habitats on Earth. The dominant bacterial groups — or *phyla*, to use the taxonomic term — are the Bacteroidota and the Bacillota (formerly known as Bacteroidetes and Firmicutes), with smaller but important contributions from Actinobacteriota, Pseudomonadota (formerly Proteobacteria), and others.

These organisms perform a staggering range of functions: fermenting dietary fibre into short-chain fatty acids that nourish the colon wall and regulate immune function, synthesising vitamins K and B12, metabolising bile acids, and — crucially — maintaining a barrier against pathogenic organisms through the ecological principle of *colonisation resistance*, a concept we will explore in depth in Chapter 4.

The skin is an entirely different world. Dry, acidic, exposed to air and ultraviolet light, it is a harsh environment by microbial standards. Yet it supports a thriving community — dominated by bacteria of the genera *Cutibacterium* (formerly *Propionibacterium*), *Staphylococcus*, and *Corynebacterium*, along with the fungus *Malassezia*, which is the most abundant eukaryotic organism on the skin surface. The composition varies dramatically between body sites: the oily terrain of the face favours lipid-loving *Cutibacterium acnes*, while the moist folds of the groin and armpit support different bacterial species altogether. Even the left and right hands of the same individual harbour measurably different communities, a finding that has intriguing implications for forensic science.

The mouth is one of the most ecologically complex sites in the body. Over 700 species of bacteria have been identified in the human oral cavity, organised into intricate biofilms — structured, surface-attached communities that were, in fact, among the first microbial ecosystems ever studied. The plaque on your teeth is a biofilm. It is also, in a very real sense, where microbiome science began: van Leeuwenhoek's original *animalcules* were scraped from between his own teeth.

The respiratory tract was, until recently, considered sterile below the larynx. It is not. The lungs harbour a low-density but distinct microbial community, and the nasopharynx — the space behind the nose — is home to species that play important roles in both health and disease. The discovery of a lung microbiome, and the ongoing debate about its significance, is one of the most active areas of current research.

The urogenital tract presents perhaps the starkest example of how the microbiome differs between sexes. The vaginal microbiome is often dominated by one or a few species of *Lactobacillus*, which produce lactic acid and maintain a low pH that inhibits pathogenic organisms. This is one of the clearest examples of mutualism in the human microbiome: the *Lactobacillus* species benefit from the nutrient-rich vaginal environment, and the host benefits from the acidic barrier they create. The urinary tract, like the lungs, was long assumed to be sterile; recent studies using sensitive molecular methods have shown that it, too, has a resident microbial community, albeit a sparse one.

And then there is **the brain**. For over a century, the central nervous system was considered the most sterile compartment in the body — a sanctuary behind the blood-brain barrier where no microorganism trod. Recent research has begun to challenge this assumption, with several groups reporting the detection of bacterial DNA and, in some cases, intact microorganisms in brain tissue from both healthy individuals and those with neurodegenerative disease. Whether these represent genuine residents, transient visitors, or contamination artefacts remains fiercely debated. We will devote an entire chapter to this question later in the book.

What emerges from this survey is a picture of extraordinary diversity and ecological sophistication. The human body is not a single environment but a continent of microclimates, each supporting its own adapted community. And these communities are not static. They change across the lifespan — from the pioneer species that colonise a newborn's gut to the diminished, less diverse communities often found in the elderly. They fluctuate with diet, with medication, with illness, with stress, with the seasons, even with the time of day.

But beneath this variation, there is also consistency. Certain core functions — the fermentation of dietary fibre, the production of certain vitamins, the maintenance of colonisation resistance — are performed reliably across individuals, even when the specific species responsible differ from person to person. This *functional redundancy* — the idea that different species can perform the same job — is a hallmark of stable ecosystems, and it gives the microbiome a resilience that pure species counts do not capture.

You Are Never Alone — And You Wouldn't Want to Be

There is one more idea to carry forward from this opening chapter, and it is perhaps the most important one.

Every organism more complex than a single cell lives in intimate association with microorganisms. This is not a quirk of human biology. It is a universal feature of life on Earth.

Corals depend on symbiotic algae for photosynthesis — without them, corals bleach and die. Leguminous plants form partnerships with nitrogen-fixing bacteria in their root nodules; without those bacteria, they cannot extract nitrogen from the atmosphere and they starve. Termites cannot digest the cellulose in wood without the microbial communities in their hindguts. Cattle and other ruminants owe their entire digestive strategy to the bacteria, archaea, fungi, and protozoa that inhabit their rumens. Aphids carry obligate bacterial symbionts, *Buchnera*, without which they cannot synthesise essential amino acids; the association is so ancient that neither organism can survive without the other. Even the mitochondria inside your own cells — the organelles that generate the energy currency of life — are the remnants of an ancient bacterial symbiosis, a merger that occurred some two billion years ago and gave rise to all complex life on Earth.

Multicellular organisms did not evolve in a sterile vacuum. They evolved *in the constant presence of microorganisms*, and they evolved to take advantage of it. The question is not whether microbes affect us. It is whether there is any aspect of our biology that they *do not* affect.

This is not the story most of us were taught. We grew up with a model of the human body as a self-contained machine — a marvel of biological engineering that happens to be occasionally invaded by hostile organisms. The microbiome tells a different story. It tells us that we are not self-contained. That we are ecosystems. That the boundary between “self” and “other” is far more porous than we imagined.

And it tells us that the organisms we have spent a century trying to eliminate, avoid, and destroy are, in many cases, the very partners we need to thrive.

In the chapters that follow, we will explore that partnership in detail — where these organisms live, how they communicate with our immune system and our brain, what happens when we damage them with antibiotics or impoverished diets, and what the latest science tells us about restoring and maintaining a healthy microbiome. Some of what we will find is reassuring. Some of it is alarming. And some of it is simply astonishing.

But it all begins with a simple shift in perspective: the recognition that we are never alone — and that this is not a problem to be solved. It is a condition to be understood.

Chapter 1 — Key References

- [REF:sender2016] Sender, R., Fuchs, S. & Milo, R. (2016). Revised estimates of the number of human and bacteria cells in the body. *PLOS Biology*, 14(8), e1002533. <https://doi.org/10.1371/journal.pbio.1002533>
- [REF:luckey1972] Luckey, T.D. (1972). Introduction to intestinal microecology. *American Journal of Clinical Nutrition*, 25(12), 1292–1294.
- [REF:berg1996] Berg, R.D. (1996). The indigenous gastrointestinal microflora. *Trends in Microbiology*, 4(11), 430–435.
- [REF:bianconi2013] Bianconi, E. et al. (2013). An estimation of the number of cells in the human body. *Annals of Human Biology*, 40(6), 463–471. <https://doi.org/10.3109/03014460.2013.807878>
- [REF:berg2020] Berg, G. et al. (2020). Microbiome definition re-visited: old concepts and new challenges. *Microbiome*, 8, 103. <https://doi.org/10.1186/s40168-020-00875-0>
- [REF:gilbert2018] Gilbert, J.A. et al. (2018). Current understanding of the human microbiome. *Nature Medicine*, 24(4), 392–400. <https://doi.org/10.1038/nm.4517>
- [REF:qin2010] Qin, J. et al. (2010). A human gut microbial gene catalogue established by metagenomic sequencing. *Nature*, 464(7285), 59–65. <https://doi.org/10.1038/nature08821>
- [REF:tierney2019] Tierney, B.T. et al. (2019). The landscape of genetic content in the gut and oral human microbiome. *Cell Host & Microbe*, 26(2), 283–295. <https://doi.org/10.1016/j.chom.2019.07.008>
- [REF:almeida2021] Almeida, A. et al. (2021). A unified catalog of 204,938 reference genomes from the human gut microbiome. *Nature Biotechnology*, 39, 105–114. <https://doi.org/10.1038/s41587-020-0603-3>
- Lederberg, J. & McCray, A.T. (2001). 'Ome sweet 'omics — a genealogical treasury of words. *The Scientist*, 15(7), 8.
- The Human Microbiome Project Consortium. (2012). Structure, function and diversity of the healthy human microbiome. *Nature*, 486(7402), 207–214.
- Strachan, D.P. (1989). Hay fever, hygiene, and household size. *BMJ*, 299(6710), 1259–1260.
- Koch, R. (1876). Untersuchungen über Bakterien: V. Die Ätiologie der Milzbrand-Krankheit. *Cohns Beiträge zur Biologie der Pflanzen*, 2(2), 277–310.
- McFall-Ngai, M. et al. (2013). Animals in a bacterial world, a new imperative for the life sciences. *Proceedings of the National Academy of Sciences*, 110(9), 3229–3236.
- Ley, R.E., Peterson, D.A. & Gordon, J.I. (2006). Ecological and evolutionary forces shaping microbial diversity in the human intestine. *Cell*, 124(4), 837–848.
- Margulis, L. (1970). *Origin of Eukaryotic Cells*. Yale University Press.

Chapter 2: Built Together — How Microbes Shaped Human Evolution

In the last chapter, we established that you are never alone — that trillions of microorganisms live in and on your body, and that most of them are partners rather than enemies. But that framing, accurate as it is, still carries a faint implication of separateness. It suggests two parties — you and them — who happen to coexist. A landlord and tenants. A host and guests.

The truth is stranger than that, and more intimate.

The microorganisms that surround you are not just living *with* you. Some of them are living *inside your genome*. Ancient viruses, embedded in your DNA millions of years ago, now perform functions so essential that you could not develop in the womb, form a placenta, fight off infections, or consolidate a memory without them. Genes borrowed from bacteria regulate processes in your cells that you think of as uniquely, irreducibly human. The very architecture of your nervous system — the synaptic machinery that allows you to read and understand this sentence — depends on a protein that was once part of a virus.

You are not a human being who happens to carry microbial passengers. You are a mosaic — a composite organism assembled, over billions of years, from parts that were never exclusively yours.

This chapter is about how that happened.

2.1 The First Great Merger: Mitochondria and the Birth of Complexity

We touched briefly on mitochondria in Chapter 1, but the story deserves more than a passing reference, because it is the founding event of everything that follows. Without it, there would be no animals, no plants, no fungi, no complex life of any kind. There would be bacteria, archaea, and nothing else.

Roughly two billion years ago — the date remains imprecise and debated, but the event itself is not — a single-celled organism, probably an archaeon, engulfed a bacterium. This was not unusual; cells consume other cells all the time. What was unusual is that, on this occasion, the engulfed bacterium was not digested. It survived inside its host. And over time, the two organisms entered into a metabolic partnership so advantageous to both that neither could survive without the other.

The engulfed bacterium, an ancestor of modern *Alphaproteobacteria*, was extraordinarily good at one thing: using oxygen to generate energy. The host cell was not. In the oxygen-rich atmosphere that was beginning to develop on the early Earth, this was a decisive advantage. The bacterium provided efficient aerobic energy production; the host provided a stable, nutrient-rich environment. Over hundreds of millions of years, the bacterium lost genes it no longer needed — genes for independent living, for building a cell wall, for motility — and transferred many of its remaining genes into the host's nuclear genome. It became an organelle: the mitochondrion.

This event — known as *endosymbiosis*, a term coined by the biologist Lynn Margulis, who championed the theory for decades against fierce resistance from the scientific establishment — was

arguably the single most consequential event in the history of life after the origin of life itself. Mitochondria gave their host cells access to vastly more energy than was available through anaerobic metabolism. That energy surplus, it is now believed, is what made biological complexity possible. Building large cells, maintaining elaborate internal structures, producing multicellular bodies — all of this is energetically expensive. Without mitochondria, the energy budget simply does not work.

Every animal cell in your body, with the exception of mature red blood cells, contains mitochondria — typically hundreds or thousands of them per cell. They retain their own small circular genome, a vestige of their bacterial ancestry, and they replicate semi-independently within the cell, dividing by a process that resembles bacterial fission far more than it resembles human cell division. They are, in the most literal sense, domesticated bacteria. Your body runs on bacterial power.

Margulis, who died in 2011, endured decades of scepticism for this idea. Her first paper on the subject was rejected by fifteen journals before being published in 1967. Today, endosymbiosis is not merely accepted; it is one of the foundational principles of cell biology. It is taught to every undergraduate biology student on Earth. And it carries a lesson that extends far beyond mitochondria: the most transformative events in evolutionary history are not always about competition. Sometimes they are about merger.

2.2 The Viral Archive: Eight Per Cent of You Is Virus

If mitochondria represent the oldest and deepest microbial contribution to human biology, the next layer is, if anything, more surprising. Scattered throughout your genome — woven into the DNA of every cell in your body — are the remnants of ancient viral infections. Not recent infections; not the kind of virus you catch and recover from. These are viruses that infected your ancestors' germ cells — their eggs or sperm — millions of years ago, inserted their genetic material into the host genome, and then, instead of killing the cell or being eliminated, became *permanent residents*.

They are called *endogenous retroviruses*, or HERVs (human endogenous retroviruses), and they are astonishingly abundant. Approximately eight per cent of the human genome — some 98,000 elements, comprising fragments, partial sequences, and a few largely intact proviral genomes — is derived from retroviruses. To put that in perspective, the protein-coding genes that we think of as “our” genes — the sequences that encode the proteins that build and run the human body — account for only about 1.5 per cent of the genome.

You are, by weight of sequence, more virus than human gene.

How did this happen? Retroviruses have a unique lifecycle. Unlike most viruses, which inject their genetic material and hijack the cell temporarily, retroviruses convert their RNA genome into DNA using an enzyme called reverse transcriptase, and then permanently integrate that DNA into the host cell's genome using another enzyme called integrase. Once integrated, the viral sequence — called a *provirus* — is replicated along with the host's own DNA every time the cell divides. If the cell happens to be a germ cell — an egg or a sperm — then the provirus will be passed to the next generation, and the next, and the next.

Most of these viral insertions are very old. Some HERV families entered the primate lineage more than 30 million years ago; others are shared with all placental mammals, placing their origin at 100 million years or more. Over such vast timescales, most endogenous retroviruses have accumulated so many mutations that they can no longer produce functional viruses. They are, in a sense, fossils —

molecular records of ancient pandemics, preserved in the genome long after the viruses themselves went extinct.

For decades, endogenous retroviruses were dismissed as “junk DNA” — genomic parasites that served no useful purpose and were simply too embedded to be removed by natural selection. This view has turned out to be profoundly wrong.

2.3 The Gift of the Placenta: Syncytins and the Viral Origins of Pregnancy

Of all the revelations to emerge from the study of endogenous retroviruses, the most striking — and the most consequential for human biology — is the discovery that the placenta, the organ that defines mammalian pregnancy, depends on genes stolen from viruses.

The key players are two proteins called Syncytin-1 and Syncytin-2. Syncytin-1 is encoded by a gene derived from the envelope protein of an endogenous retrovirus called HERV-W. Syncytin-2 comes from a different retrovirus, HERV-FRD. Both were identified in the early 2000s, and both turn out to be essential for the formation of the placental structure that nourishes the developing foetus.

To understand why, you need to know a little about how the placenta works. The outer surface of the human placenta — the interface between maternal and foetal tissue — is composed of a continuous, multinucleated layer called the *syncytiotrophoblast*. This is not a sheet of individual cells sitting side by side. It is a single, vast, fused structure — a syncytium — in which the cell membranes between adjacent cells have dissolved, allowing their contents to merge into one continuous mass. This syncytium is critical. It provides the enormous surface area needed for nutrient and gas exchange between mother and foetus, and it acts as an immunological barrier that prevents the mother’s immune system from rejecting the foetus, which is, after all, genetically half foreign.

How does this syncytium form? By cell-cell fusion — a process in which the membranes of adjacent cells merge. And this is precisely what retroviral envelope proteins evolved to do. In an active retrovirus, the envelope protein’s job is to fuse the viral membrane with the host cell membrane, allowing the virus to enter the cell. The syncytins have been repurposed: instead of fusing a virus to a cell, they fuse placental cells to each other, creating the syncytiotrophoblast.

The discovery was made by teams working independently in France and the United States. In 2000, the group led by Thierry Heidmann identified Syncytin-1 and showed that it was expressed at high levels in the placenta and could mediate cell-cell fusion. Subsequent work by several groups confirmed that Syncytin-2, from a different retroviral source, plays a complementary role and additionally has immunosuppressive properties — dampening the maternal immune response at the foetal-maternal interface.

The implications are extraordinary. Without these co-opted viral genes, the placental syncytiotrophoblast could not form, and mammalian pregnancy as we know it could not occur. Knock out syncytin genes in mice (which have their own, independently captured syncytins — a point we will return to) and the placentas are severely defective, leading to embryonic death. The virus that infected our ancestors did not merely leave a trace in the genome. It provided the molecular machinery for one of the defining innovations of mammalian life.

And the story goes deeper — much deeper.

Remarkably, syncytin-like genes have been captured independently by different mammalian lin-

eages from *different* retroviruses on at least seven separate occasions across evolutionary history. Rodents use different syncytins from primates, which use different syncytins from rabbits, which use different syncytins from carnivores. Even the opossums — marsupials, which diverged from eutherian mammals some 148 million years ago — carry their own captured syncytin, called Syncytin-Opo1. The same evolutionary solution — co-opting a retroviral envelope protein for placental fusion — has been reinvented again and again. This is a powerful argument that the co-option is not accidental. Natural selection has actively favoured retaining and repurposing these viral genes because the function they provide is so valuable that evolution keeps rediscovering it.

But the most astonishing piece of evidence does not come from mammals at all. It comes from a lizard.

The *Mabuya* skinks are a genus of small, smooth-scaled lizards found in South America and Africa. Most lizards lay eggs. The *Mabuya* does not. Sometime around 25 million years ago, this lineage made the transition from oviparity — egg-laying — to viviparity: live birth, with embryos nourished not by a yolk sac but by a placenta-like structure that is, in its architecture, strikingly similar to the placenta of eutherian mammals. *Mabuya* embryos develop inside the mother, receiving nearly all of their nutrition across a materno-foetal interface that includes a syncytialised cell layer. The eggs are virtually yolkless. The dependence on maternal provisioning is almost total.

In 2017, a team led by Guillaume Cornelis and Thierry Heidmann — the same Heidmann who had identified human Syncytin-1 seventeen years earlier — examined placental tissue from *Mabuya* skinks. They found exactly what you might now predict: a retroviral envelope gene, which they named Syncytin-Mab1, that was expressed at high levels in the placenta, that could drive cell-cell fusion, and that had been conserved across all *Mabuya* species tested, spanning 25 million years of evolution. They also identified its receptor, a transmembrane protein called MPZL1. It was, in every functional respect, a bona fide syncytin — captured from an entirely different retrovirus than any of the mammalian syncytins, in a lineage separated from mammals by more than 300 million years of evolution.

A lizard, descended from egg-laying ancestors, had independently evolved a mammalian-style placenta using the same molecular trick: enslaving a retroviral envelope protein to fuse cells at the materno-foetal interface. This is convergent evolution at its most extraordinary — not just convergence of form, but convergence of mechanism. Two vastly distant lineages, faced with the same biological challenge (how to nourish an embryo internally), arrived at the same solution (co-opt a virus).

The *Mabuya* story also illuminates a broader evolutionary spectrum that is worth pausing to appreciate. Not all placentas are alike, and the variation tells us something about the depth and duration of viral co-option.

At one end of the spectrum sit the *monotremes* — the platypus and the echidnas — which lay eggs. They diverged from other mammals roughly 166 million years ago, before the evolution of placentation. They have no syncytins. They have no placenta. They nourish their young with milk, secreted through patches of skin rather than through nipples, but embryonic development occurs entirely within the egg.

Next come the *marsupials* — kangaroos, koalas, opossums, and their kin. Marsupials do have a placenta, but it is a relatively simple structure: a *choriovitelline* placenta, derived primarily from the yolk sac. Gestation is short. A newborn kangaroo joey is born after just 30-odd days, weighing less than a gram — essentially a foetus that completes its development externally, attached to a

nipple inside the mother's pouch. The marsupial placenta provides nourishment for a brief period, but the heavy lifting of development happens after birth. Marsupials have their own captured syncytin (Syncytin-Opo1 in opossums), but the lesser complexity of their placental structure suggests a shallower integration of retroviral machinery.

Then come the *eutherian* mammals — the group that includes humans, along with rodents, carnivores, bats, whales, and most of the mammals you can readily name. Eutherians possess a *chorioallantoic* placenta, a far more elaborate structure that enables extended gestation, deep embryonic implantation, and intimate exchange between maternal and foetal blood supplies. A human pregnancy lasts nine months. An elephant's lasts 22 months. This prolonged internal development — and the larger, more mature offspring it produces — would not be possible without the sophisticated placental architecture that syncytins help to build.

The pattern is suggestive: the more thoroughly a lineage has co-opted retroviral envelope genes for placental function, the more elaborate its placenta, the longer its gestation, and the more developed its offspring at birth. The transition from egg to pouch to womb is not a single evolutionary leap. It is a graduated sequence, and at each stage, viruses have been recruited to enable the next step.

And now we know that this is not even a uniquely mammalian phenomenon. The Mabuya lizard tells us that the transition from oviparity to viviparity — one of the most profound shifts in reproductive strategy that any vertebrate lineage can undergo — can be driven by the same retroviral co-option event in a reptile. The virus does not care whether its host is a mammal or a lizard. It provides the molecular tool. Evolution does the rest.

2.4 How a Virus Wired the Nervous System: RetroMyelin and the Insulation of Thought

If syncytins show how viruses enabled mammalian reproduction, the next story shows how they enabled complex nervous systems — and it comes from a paper published in *Cell* as recently as February 2024.

Myelin is the fatty insulating sheath that wraps around nerve fibres — axons — in both the brain and spinal cord. If you have ever seen a diagram of a nerve cell, you will recognise myelin as the series of sausage-shaped segments along the axon, separated by small gaps called nodes of Ranvier. This arrangement is not merely protective; it is the basis of a conduction strategy called *saltatory conduction*, in which electrical impulses leap from one gap to the next rather than travelling continuously along the fibre. The result is a dramatic increase in both the speed and energy efficiency of nerve transmission.

How dramatic? Consider what the nervous system would look like without it. In unmyelinated nerve fibres, conduction speed is roughly proportional to the diameter of the fibre: to transmit faster, you need a fatter axon. The squid giant axon, a classic object of neurophysiology, achieves rapid escape responses by brute force — it is nearly a millimetre in diameter, enormous by neural standards. With myelination, a vertebrate axon only twenty micrometres across — fifty times thinner — can conduct impulses at the same speed. The diameter ratio translates to cross-sectional area: an unmyelinated axon needs roughly 2,500 times the cross-section to match the conduction velocity of its myelinated counterpart.

Now scale that up to the human spinal cord, which contains millions of axons bundled together.

Estimates vary, but if the spinal cord were composed entirely of unmyelinated fibres and needed to carry the same information at the same speed, it would need to be vastly thicker — by some calculations, the cross-section would have to increase by a factor of ten or more. A spinal cord the diameter of a garden hose would become something closer to a drainpipe. The human body plan, with its narrow vertebral column and slender neck, would simply not be viable. Compact myelination does not merely make nervous systems faster. It makes complex nervous systems *physically possible*.

Myelin first appears in the fossil record and in living species with the emergence of jawed vertebrates — the gnathostomes — roughly 450 million years ago. Sharks, rays, and all bony fishes have myelin. Lampreys and hagfish — the jawless vertebrates, or agnathans — do not. They have glial cells that wrap loosely around axons but do not form the tight, compacted myelin sheath that enables saltatory conduction. The appearance of true myelin coincides with the explosive diversification of vertebrate body plans and neural complexity that characterises the gnathostome radiation. This is not a coincidence.

But what molecular event triggered the emergence of myelination? Despite decades of research, the answer remained elusive until Tanay Ghosh and colleagues at the University of Cambridge identified a startling piece of the puzzle: a retrotransposon of retroviral origin, which they named *RetroMyelin*.

RetroMyelin is derived from an endogenous retroviral sequence called RNLTR12-int, belonging to the ERV1 class of retrotransposons. Ghosh's team showed that this sequence does not encode a protein. Instead, it produces an RNA molecule that binds directly to the transcription factor SOX10 — a master regulator of oligodendrocyte development — and is essential for switching on the gene for *myelin basic protein* (MBP), the major structural component of the myelin sheath. Without RetroMyelin RNA, MBP is not expressed at normal levels, and myelination fails.

The team demonstrated this in rodent oligodendrocytes, then validated it in zebrafish and frogs — two vertebrate classes separated by hundreds of millions of years of evolution. RetroMyelin-like sequences were found in the genomes of all jawed vertebrates examined, but not in lampreys or hagfish. The implication is extraordinary: the retroviral sequence was endogenised into the genome of an ancestral jawed vertebrate, and its RNA was co-opted to regulate the gene that makes myelin possible. Without this viral contribution, the myelin sheath — and everything it enables — may never have evolved.

Pause for a moment to consider what that means. Your ability to move, to see, to think at speed — all of this depends on myelinated axons conducting impulses in milliseconds rather than the sluggish conduction that unmyelinated fibres would provide. The entire architecture of the vertebrate central nervous system is built on the assumption that myelin is available. And myelin is available because a virus, hundreds of millions of years ago, inserted its genetic material into the genome of a fish-like ancestor, and that material was repurposed into a regulatory switch for one of the most consequential genes in neural biology.

The octopus offers a poignant counterpoint. Cephalopods are, by any measure, the most intelligent invertebrates — capable of problem-solving, tool use, and behaviours that suggest something approaching cognition. Their nervous system is extraordinarily sophisticated, with roughly 500 million neurons (comparable to a dog) distributed between a central brain and semi-autonomous ganglia in each arm. But octopuses do not have myelin. Their axons are unmyelinated. And this imposes a fundamental constraint on neural architecture: to achieve adequate conduction speeds, their nerve fibres must be relatively thick, limiting how many can be packed into a given space. The

octopus has pushed unmyelinated neural design to remarkable heights, but there is a ceiling — a physical limit to how complex an unmyelinated nervous system can become before it runs out of room.

It is tempting to speculate — and this is speculation, to be clear — that the lack of myelin is one reason cephalopod intelligence, impressive as it is, has not scaled to the levels seen in vertebrates. The octopus brain is a monument to what evolution can achieve without viral co-option. The vertebrate brain is a monument to what becomes possible with it.

2.5 How a Virus Gave Us Memory: The ARC Protein

If the placenta story is about how viruses helped build the body, and the myelin story about how they wired the nervous system, the ARC story is about how they enabled the mind's most intimate function: the ability to remember.

The ARC protein (Activity-Regulated Cytoskeleton-associated protein, also known as Arg3.1) has been known to neuroscientists for decades as a key player in *synaptic plasticity* — the process by which the connections between neurons are strengthened or weakened in response to experience. Synaptic plasticity is the molecular basis of learning and memory. Without it, the brain cannot encode new information, consolidate short-term memories into long-term ones, or adapt its circuitry to a changing world.

What makes ARC unusual is its structure. In 2018, two research groups — one led by Jason Bhatt and another by Elissa Bhatt and colleagues at the University of Utah, and a parallel effort by Ed Bhatt and colleagues — published papers in the journal *Cell* revealing something unexpected: ARC protein spontaneously assembles into structures that closely resemble viral capsids. A capsid is the protein shell that packages and protects a virus's genetic material. ARC, it turned out, was forming capsid-like particles, encapsulating its own messenger RNA inside them, and transferring those RNA-containing capsids *from one neuron to another* across the synaptic gap.

In other words, neurons were communicating using virus-like particles. This was not a metaphor. It was literally viral biology, repurposed for neural signalling.

The reason became clear when researchers examined ARC's evolutionary origin. The gene encoding ARC is derived from the *gag* gene of an ancient retrotransposon — a mobile genetic element closely related to retroviruses. The *gag* gene in a retrovirus encodes the structural proteins of the viral capsid. In ARC, this same capsid-forming capability has been co-opted to create a system for intercellular RNA transfer in the nervous system.

The functional importance is not in doubt. Mice lacking the ARC gene can learn simple tasks but cannot consolidate long-term memories. They can acquire new information but cannot retain it. ARC is expressed rapidly and transiently in neurons that are actively engaged in learning, and its expression is essential for the synaptic changes that convert fleeting experience into durable memory. The discovery that this crucial protein is, at its molecular heart, a domesticated viral gene adds a remarkable layer to our understanding of what the brain is and how it came to be.

There is something philosophically arresting about this. The capacity for memory — the ability to carry your past with you, to learn from experience, to build an identity across time — is one of the things we consider most fundamentally human. And it depends on molecular machinery borrowed from a virus. The self that remembers is, in part, a viral construction.

2.6 How a Transposon Built the Immune System: RAG and V(D)J Recombination

If the ARC story connects viruses to neural function, the story of RAG connects mobile genetic elements to the immune system — and the stakes could hardly be higher.

Your adaptive immune system — the branch of immunity that produces antibodies, that generates memory T cells after a vaccination, that distinguishes the influenza virus you encountered last winter from the one circulating this year — has a problem to solve that is, on its face, impossible. It must be prepared to recognise virtually any molecular structure it might encounter: every virus, every bacterium, every toxin, every aberrant protein on the surface of a cancer cell. The number of potential threats is effectively infinite. Yet the human genome, large as it is, contains only about 20,000 protein-coding genes. How do you generate an effectively infinite diversity of immune receptors from a finite genome?

The answer is a process called *V(D)J recombination*, and it is one of the most elegant solutions in all of biology. Rather than encoding a separate gene for every possible antibody or T cell receptor, the immune system stores the genes in fragments — variable (V), diversity (D), and joining (J) segments — and then physically rearranges them during the development of each new immune cell, cutting and splicing the DNA to produce a unique combination. Because the number of possible combinations is astronomical, each B cell and T cell ends up with a receptor that is essentially unique, never before seen in evolutionary history. Your immune system does not inherit its specificity. It *generates* it, on the fly, by rewriting its own DNA.

The enzymes that perform this rearrangement are called RAG1 and RAG2 (Recombination-Activating Genes 1 and 2). Without them, V(D)J recombination does not occur. Individuals who are born with mutations that inactivate either RAG gene have *severe combined immunodeficiency* — SCID — sometimes called “bubble boy disease.” They have no functional adaptive immune system. No antibodies. No antigen-specific T cells. Without medical intervention, even trivial infections are fatal.

So where did RAG1 and RAG2 come from?

The molecular structure of the RAG1/RAG2 complex is strikingly similar to that of a *transposase* — the enzyme used by transposons (mobile genetic elements, sometimes called “jumping genes”) to cut themselves out of one location in the genome and paste themselves into another. In 2005, Vladimir Kapitonov and Jerzy Jurka identified a class of DNA transposons in the genomes of various organisms, which they named *Transib* transposons, that shared unmistakable structural and sequence similarity with RAG1. Subsequent work demonstrated that the RAG1/RAG2 complex can, in experimental conditions, function as a transposase — cutting and pasting DNA in the same manner as its ancestral transposon.

The conclusion is now widely accepted: the RAG recombinase that drives V(D)J recombination — and thus the entire adaptive immune system — is derived from a transposon that invaded the genome of an early vertebrate ancestor roughly 500 million years ago. What was once a selfish genetic parasite, hopping through the genome for its own replication, was captured, tamed, and repurposed into the molecular engine of adaptive immunity.

Without this event, there would be no antibodies. No vaccination. No immunological memory. No

ability to mount a targeted immune response against a specific pathogen. The entire edifice of adaptive immunity — the system that protects you from infection, that enables organ transplantation, that forms the basis of cancer immunotherapy — rests on a foundation built by a transposon.

2.7 Not Junk: Transposable Elements as an Evolutionary Toolkit

The stories of syncytins, ARC, and RAG are dramatic, but they are not isolated cases. They are the most visible examples of a far broader phenomenon: the pervasive contribution of mobile genetic elements — transposons, retrotransposons, and endogenous retroviruses — to the evolution of the human genome.

Roughly 45 per cent of the human genome is composed of recognisable transposable element sequences. Combined with the approximately eight per cent derived from endogenous retroviruses, this means that well over half of your DNA has its origins in mobile genetic elements — “parasitic” sequences that were once dismissed as junk but are increasingly recognised as a vast evolutionary toolkit.

The mechanisms by which these elements contribute to genome evolution are diverse and, in many cases, subtle. Transposable elements can alter gene expression by inserting themselves near existing genes and providing new regulatory sequences — promoters, enhancers, insulators, and splice sites that change when, where, and how much a gene is expressed. In some cases, these regulatory changes have been profoundly consequential.

Consider a few examples. The *MER20* element, a DNA transposon, inserted itself near genes involved in the hormonal regulation of pregnancy in the ancestor of placental mammals. Its presence rewired those genes to respond to progesterone, contributing to the evolution of the decidual cell — a cell type essential for embryo implantation in the uterine wall. Without this transposon-derived regulatory sequence, the signalling cascade that prepares the uterus for pregnancy may not have evolved as it did.

LINE-1 retrotransposons, the most abundant mobile elements in the human genome, continue to be active today. They are now known to play a role in brain development: LINE-1 elements are actively transcribed and undergo new insertions specifically in neural progenitor cells, generating genetic diversity between individual neurons. This somatic mosaicism — each neuron in your brain carries a slightly different genome — is thought to contribute to the extraordinary functional diversity of the nervous system. The idea that your brain’s complexity is partly generated by ongoing retrotransposon activity is still being explored, but the evidence is accumulating steadily.

And at the regulatory level, endogenous retroviruses have contributed thousands of transcription factor binding sites to the human genome. A study by Edward Chuong, Nels Elde, and Cédric Feschotte, published in *Science* in 2016, demonstrated that a specific family of endogenous retroviruses — MER41 — provided binding sites for the interferon-gamma signalling pathway, a critical component of the antiviral immune response. Deleting these HERV-derived binding sites in cell lines abolished the interferon-driven expression of certain immune genes. In other words, part of your antiviral defence system is regulated by sequences donated by viruses.

There is an irony here that is almost too neat. The immune system’s ability to fight viruses depends, in part, on genetic material provided by viruses. The deep evolutionary history of the genome is not

a story of hosts and parasites locked in eternal combat. It is a story of combat, détente, domestication, and — eventually — collaboration.

2.8 Bacterial Genes in the Human Genome

Viruses and transposons are not the only microbial contributors to the human genome. There is growing evidence that genes have been transferred directly from bacteria into the genomes of animals, including our own ancestors, through a process called *horizontal gene transfer* (HGT).

Horizontal gene transfer is commonplace among bacteria. It is, in fact, the principal mechanism by which antibiotic resistance spreads through bacterial populations — a topic we will explore in detail in later chapters. But for a long time, HGT was assumed to be largely confined to the prokaryotic world. The conventional view held that animal genomes acquired new genes primarily through mutation, duplication, and recombination of existing sequences, not by importing genes wholesale from other organisms.

That view has been challenged. Studies of the human genome and those of other animals have identified dozens to hundreds of genes that appear to have been acquired by horizontal transfer from bacteria or other microorganisms at various points in animal evolutionary history. The numbers and specific genes identified vary between studies, and there is legitimate scientific debate about how many of these apparent transfers are genuine versus artefacts of contamination or analytical error. But the phenomenon itself is now broadly accepted, even if its quantitative importance remains uncertain.

Among the more convincing cases are genes involved in metabolism — enzymes that perform biochemical reactions not typically found in the ancestral eukaryotic repertoire. The hyaluronan synthase gene in vertebrates, for example, has been proposed to have a bacterial origin, though this remains debated. What is less debated is that horizontal gene transfer from bacteria to eukaryotes occurs regularly in organisms that live in intimate association with bacterial symbionts — insects, nematodes, and other invertebrates — and that it has contributed functional genes to their hosts.

The broader point is that the boundaries between kingdoms of life are more porous than textbook diagrams suggest. DNA flows between organisms — between bacteria, between viruses and hosts, and, at least occasionally, between bacteria and the animal lineages that eventually gave rise to us. The genome is not a sealed vault. It is an open ledger, and many of the entries were written by other organisms.

2.9 Evolution as Collaboration

Step back from the details for a moment and consider the larger picture that this chapter paints.

Mitochondria: captured bacteria that provide the energy for all complex life. Endogenous retroviruses: viral sequences that now encode the proteins required for placental formation, without which no mammal — including us — would exist. RetroMyelin: a retroviral sequence that controls myelin production, without which fast neural transmission and compact brains are impossible. ARC: a domesticated viral gene that enables memory and learning. RAG: a tamed transposon that built

the adaptive immune system. Transposable elements: genomic wanderers that have been repurposed into regulatory switches, neural diversity generators, and immune signalling components. Bacterial genes: direct horizontal contributions to the animal genetic repertoire.

None of these are peripheral. They are not curiosities or footnotes. They are central to what it means to be a complex, multicellular, thinking, remembering, reproducing organism. Remove any one of these contributions and the human body does not function. Remove all of them and there is no human body.

The traditional narrative of evolution — the one most of us learned in school — is a story of competition. Survival of the fittest. Nature red in tooth and claw. Organisms competing for resources, with the winners passing on their genes and the losers going extinct. This narrative is not wrong. Competition is a real and powerful evolutionary force. But it is dangerously incomplete.

The history of life on Earth is also, and perhaps more profoundly, a history of collaboration: of organisms merging, co-opting, and repurposing each other's biological innovations. The mitochondrial endosymbiosis was not a competitive victory. It was a merger. The co-option of syncytins was not an arms race. It was biological recycling. The domestication of RAG was not conquest. It was the transformation of a parasite into a tool.

Lynn Margulis, the biologist who championed the endosymbiotic origin of mitochondria, saw this clearly. She argued, against the prevailing Darwinian orthodoxy of her time, that symbiosis — not just competition — was a major driver of evolutionary innovation. “Life did not take over the globe by combat,” she wrote, “but by networking.” Her view was considered radical in the 1960s. It is now mainstream, supported by a weight of genomic evidence that she could not have imagined.

This perspective has implications that extend well beyond evolutionary biology. If we have been *built* by microorganisms — if our deepest biological innovations are borrowed, co-opted, and repurposed from viruses, bacteria, and mobile genetic elements — then the adversarial framework we inherited from germ theory is not just incomplete. It is, at a fundamental level, incoherent. We cannot declare war on the microbial world without declaring war on ourselves, because the microbial world is inside us, woven into the most essential parts of what we are.

This does not mean that all microorganisms are benign. It does not mean that infectious diseases are not real or that antibiotics are not sometimes necessary. It means that the relationship between humans and microbes is not a war. It is an ancient, intricate, evolving partnership — one that has shaped every aspect of our biology, from the organelles in our cells to the architecture of our brains. We are, in the most literal and profound sense, built together.

In the next chapter, we will see how that partnership is renewed in each human life — how you, as an individual, acquired the microbial community that makes you who you are.

Chapter 2 — Key References

[To be populated with full citations during manuscript preparation]

- Margulis, L. (1967). On the origin of mitosing cells. *Journal of Theoretical Biology*, 14(3), 225–274.
- Sender, R., Fuchs, S. & Milo, R. (2016). Revised estimates for the number of human and bacteria cells in the body. *Cell*, 164(3), 337–340.

- Lander, E.S. et al. (2001). Initial sequencing and analysis of the human genome. *Nature*, 409(6822), 860–921. [~8% HERV content]
- Mi, S. et al. (2000). Syncytin is a captive retroviral envelope protein involved in human placental morphogenesis. *Nature*, 403(6771), 785–789.
- Blond, J.L. et al. (2000). An envelope glycoprotein of the human endogenous retrovirus HERV-W is expressed in the human placenta and fuses cells expressing the type D mammalian retrovirus receptor. *Journal of Virology*, 74(7), 3321–3329.
- Dupressoir, A. et al. (2009). Syncytin-A and syncytin-B, two fusogenic placenta-specific murine envelope genes of retroviral origin conserved in Muridae. *Proceedings of the National Academy of Sciences*, 106(29), 12127–12132.
- Cornelis, G. et al. (2017). An endogenous retroviral envelope syncytin and its cognate receptor identified in the viviparous placental Mabuya lizard. *Proceedings of the National Academy of Sciences*, 114(51), E10991–E11000.
- Denner, J. (2017). Function of a retroviral envelope protein in the placenta of a viviparous lizard. *Proceedings of the National Academy of Sciences*, 114(51), 13315–13317.
- Cornelis, G. et al. (2015). Retroviral envelope gene captures and syncytin exaptation for placentation in marsupials. *Proceedings of the National Academy of Sciences*, 112(5), E487–E496.
- Lavialle, C. et al. (2013). Paleovirology of ‘syncytins’, retroviral env genes exapted for a role in placentation. *Philosophical Transactions of the Royal Society B*, 368(1626), 20120507.
- Ghosh, T. et al. (2024). A retroviral link to vertebrate myelination through retrotransposon-RNA-mediated control of myelin gene expression. *Cell*, 187(4), 814–830.
- Ghosh, T. et al. (2024). A retroviral link to vertebrate myelination through retrotransposon-RNA-mediated control of myelin gene expression. *Cell*, 187(4), 964–981.
- Pastuzyn, E.D. et al. (2018). The neuronal gene Arc encodes a repurposed retrotransposon Gag protein that mediates intercellular RNA transfer. *Cell*, 172(1–2), 275–288.
- Ashley, J. et al. (2018). Retrovirus-like Gag protein Arc1 binds RNA and traffics across synaptic boutons. *Cell*, 172(1–2), 262–274.
- Kapitonov, V.V. & Jurka, J. (2005). RAG1 core and V(D)J recombination signal sequences were derived from Transib transposons. *PLoS Biology*, 3(6), e181.
- Huang, S. et al. (2016). Discovery of an active RAG transposon illuminates the origins of V(D)J recombination. *Cell*, 166(2), 468–481.
- Chuong, E.B., Elde, N.C. & Feschotte, C. (2016). Regulatory evolution of innate immunity through co-option of endogenous retroviruses. *Science*, 351(6277), 1083–1087.
- Feschotte, C. & Gilbert, C. (2012). Endogenous viruses: insights into viral evolution and impact on host biology. *Nature Reviews Genetics*, 13(4), 283–296.
- Muotri, A.R. et al. (2005). Somatic mosaicism in neuronal precursor cells mediated by L1 retrotransposition. *Nature*, 435(7044), 903–910.
- Lynch, V.J. et al. (2015). Ancient transposable elements transformed the uterine regulatory landscape and transcriptome during the evolution of mammalian pregnancy. *Cell Reports*, 10(4), 551–561.

Chapter 3: Your First Inheritance — Colonisation Across a Lifetime

The last chapter told the story of evolution — how, across billions of years, microorganisms wove themselves into the genetic fabric of complex life. The placenta, the myelin sheath, the immune system, the architecture of memory — all of these, we saw, bear the fingerprints of ancient microbial partnerships.

But that is the story of the species. There is also the story of you.

Every human being who has ever lived began as a near-sterile collection of cells. The fertilised egg that would become you contained a human genome, mitochondria (those ancient bacterial endosymbionts), and the retroviral sequences embedded in your DNA. But it did not contain a microbiome. That had to be acquired. And the way you acquired it — the particular sequence of microbial encounters that seeded, shaped, and matured your personal microbial community — turns out to have consequences that can last a lifetime.

This chapter is about that process: how you inherited not just genes from your parents, but microbes. How those microbes arrived, in what order, and why it matters. And how the community they formed has changed — and will continue to change — at every stage of your life, from the moment of your birth to the final years of your old age.

3.1 Before Birth: The Sterile Womb Debate

For most of the history of medicine, the womb was considered sterile. The foetus, protected behind the membranes of the amniotic sac, was thought to develop in a microbe-free environment, encountering its first microorganisms only during or after birth. This idea — the *sterile womb paradigm* — was deeply embedded in obstetric practice and neonatal medicine. It seemed to make intuitive sense: the placenta acts as a barrier, the amniotic fluid is normally clear, and intrauterine infection, when it occurs, is treated as a pathological event.

In the 2010s, a series of studies challenged this view. Researchers reported detecting bacterial DNA in the placenta, in amniotic fluid, and in meconium — the first stool of a newborn, formed before birth. These findings generated considerable excitement. If the womb was not sterile, then microbial colonisation might begin before birth, and the mother might be seeding her child's microbiome in utero through the placenta or the bloodstream. The implications for understanding early immune development were potentially enormous.

The excitement was premature. Subsequent work, using more rigorous controls for contamination, cast serious doubt on many of the initial findings. The problem is a technical one that bedevils all low-biomass microbiome studies: when you are looking for tiny quantities of bacterial DNA in a sample, the reagents you use to extract and amplify that DNA may themselves contain trace amounts of bacterial DNA. This *kit contamination* can produce results that look like a genuine microbial community but are in fact an artefact of the laboratory process. Several large, carefully

controlled studies published between 2019 and 2023 found that when contamination controls were rigorously applied, the placenta and amniotic fluid of healthy pregnancies showed no convincing evidence of a resident microbial community.

The current scientific consensus — though “consensus” is perhaps too strong a word for a field that remains actively contested — leans toward the traditional view: the healthy womb is, for practical purposes, sterile or very nearly so. There may be transient microbial exposures during pregnancy, particularly in the third trimester, but there is no established evidence of a sustained, resident microbiome in the uterus or placenta during normal pregnancy.

This matters because it means that the main event — the founding of your microbiome — happens at birth. And the manner of your birth has a great deal to do with what that founding community looks like.

3.2 The Birth Canal: Your First Microbial Baptism

A vaginal birth is, from a microbiological perspective, a controlled inoculation.

As the infant passes through the birth canal, it is bathed in the mother’s vaginal secretions — a fluid rich in microorganisms, dominated in most women by one or a few species of *Lactobacillus*. The baby’s skin, mouth, and nasal passages are coated with these organisms. It swallows them. It inhales them. In the hours that follow, the baby is further exposed to microorganisms from the mother’s skin, her gut (through the inevitable faecal contamination that accompanies most vaginal deliveries — a fact that is rarely mentioned in antenatal classes but is microbiologically important), and her breast. Within minutes of birth, the infant’s body is no longer sterile. The colonisation has begun.

The first organisms to arrive are called *pioneer species*, borrowing a term from ecology. Just as the first plants to colonise a bare lava flow change the soil chemistry and create conditions for later arrivals, the pioneer bacteria that colonise a newborn’s gut alter its environment in ways that favour or disfavour subsequent colonisers. They consume the available oxygen, driving down the oxygen concentration in the gut lumen and creating the anaerobic conditions that the dominant adult gut bacteria — the obligate anaerobes of the phyla Bacteroidota and Bacillota — require to thrive. They interact with the nascent immune system, beginning the long process of education that will teach the infant’s immune cells to distinguish friend from foe. They occupy ecological niches, establishing the first lines of colonisation resistance.

The identity of these pioneers matters. Studies comparing the microbiomes of vaginally delivered infants with those born by caesarean section have consistently found significant differences in the composition of the early gut microbiota. Vaginally delivered babies tend to be colonised first by *Lactobacillus*, *Prevotella*, and *Bacteroides* species — organisms derived from the mother’s vaginal and intestinal microbiota. Caesarean-delivered babies, who bypass the birth canal entirely, are instead colonised first by organisms from the hospital environment and the mother’s skin: *Staphylococcus*, *Corynebacterium*, *Cutibacterium*, and sometimes opportunistic organisms such as *Klebsiella* and *Enterococcus*.

These differences are real, well-documented, and consistently replicated across studies. But their long-term significance remains the subject of genuine scientific debate — and it is important to be honest about this, because the topic has generated more anxiety among expectant parents than

almost any other finding in microbiome science.

The concern is that caesarean-born infants, by missing the vaginal inoculation, may develop a less optimal microbiome that predisposes them to immune-related diseases later in life. Epidemiological studies have indeed found statistical associations between caesarean delivery and modestly increased rates of asthma, allergic disease, type 1 diabetes, and obesity in childhood. These associations are consistent across multiple large cohorts and have been replicated internationally.

However — and this is a critical “however” — association is not causation. Caesarean deliveries are not performed at random. They are more common in mothers with gestational diabetes, pre-eclampsia, and obesity; in pregnancies complicated by foetal distress; and in mothers who have had previous caesarean deliveries. These underlying conditions are themselves associated with altered metabolic and immune profiles in offspring. Disentangling the effect of the delivery mode from the effects of the conditions that prompted it is fiendishly difficult.

Moreover, the differences in microbiome composition between vaginally and caesarean-delivered infants, while real in the first weeks and months, tend to diminish substantially by six to twelve months of age. By the time a child is eating solid food and exploring the world on hands and knees, the microbiome has been so thoroughly reshaped by diet, environment, and host genetics that the mode of delivery is no longer the dominant influence.

This is not to say that birth mode does not matter. It may well matter, particularly for infants born by caesarean who are also formula-fed and exposed to early antibiotics — a combination that removes multiple sources of microbial inoculation simultaneously. But the popular narrative that caesarean birth causes lasting immune damage through a deficient microbiome is, at best, an oversimplification of a far more nuanced picture.

One curious practice that has emerged from this concern deserves mention: *vaginal seeding*, in which a gauze swab is placed in the mother’s vagina before a planned caesarean delivery and then wiped across the newborn’s face, mouth, and body immediately after birth, in an attempt to transfer vaginal microorganisms to the baby. The practice has gained a following among some parents and birth practitioners, but professional medical organisations have generally advised against it. The evidence that it alters the infant microbiome in a meaningful or lasting way is thin, and there is a non-trivial risk of transferring pathogenic organisms — including Group B *Streptococcus* and herpes simplex virus — that the vaginal canal may harbour. It is, for now, an intervention in search of evidence rather than an evidence-based intervention.

3.3 Breastmilk: The Microbial Care Package

If birth is the first inoculation, breastfeeding is the first act of targeted ecological management.

Human breastmilk is not merely food. It is a biological system of extraordinary sophistication, and its relationship with the infant microbiome is one of the most elegant examples of co-evolved mutualism in human biology.

The most striking feature, from a microbiome perspective, is the presence of *human milk oligosaccharides* — HMOs. These are complex sugars, and human milk contains an astonishing diversity of them: more than 200 structurally distinct HMOs have been identified, making them collectively the third most abundant solid component of breastmilk, after lactose and fat. An exclusively breastfed infant consumes roughly 10 to 15 grams of HMOs per day.

Here is the remarkable thing: the infant cannot digest HMOs. Human cells do not produce the enzymes needed to break them down. The baby's own digestive system derives no nutritional benefit from them whatsoever. So why does the mother expend considerable metabolic energy — synthesising 10 grams a day of complex carbohydrates is not cheap — producing sugars that her baby cannot use?

Because the baby is not the intended consumer. The bacteria are.

HMOs are selectively metabolised by specific species of *Bifidobacterium* — particularly *Bifidobacterium longum* subspecies *infantis* (*B. infantis*), an organism exquisitely adapted to life in the breastfed infant gut. *B. infantis* possesses a suite of enzymes — notably including specific glycosidases and transport proteins — that allow it to import and digest HMOs with extraordinary efficiency. Most other gut bacteria cannot do this. The result is that HMOs act as a highly selective prebiotic: they feed *Bifidobacterium* and starve almost everything else.

The consequences cascade through the infant gut ecosystem. *Bifidobacterium* species, fuelled by HMOs, proliferate rapidly and come to dominate the gut microbiota of breastfed infants, often accounting for 60 to 90 per cent of the total bacterial community. As they ferment HMOs, they produce short-chain fatty acids — particularly acetate and lactate — that lower the pH of the gut lumen. This acidic environment further inhibits the growth of many potential pathogens while nourishing the cells of the infant colon wall. The *Bifidobacterium* also interact directly with the developing immune system, promoting the maturation of regulatory T cells and the establishment of immune tolerance.

The mother, in other words, is not just feeding her baby. She is farming her baby's microbiome. She is producing a sophisticated prebiotic that selectively cultivates the organisms she wants her infant to carry, while suppressing the organisms she does not. And she is doing this through her own breast, using metabolic resources that could otherwise be directed to her own nutrition.

This has implications for infant formula. Modern formulas are nutritionally sophisticated — they can sustain healthy growth and development — but until very recently, they did not contain HMOs. Without HMOs, the selective pressure that drives *Bifidobacterium* dominance is absent, and formula-fed infants typically develop a more diverse but less *Bifidobacterium*-dominated gut community, with higher proportions of *Clostridium*, *Enterobacteriaceae*, and *Bacteroides*. Whether this difference has lasting health consequences remains debated, but it is one reason that breastfeeding advocacy organisations emphasise the microbiome benefits of breastfeeding alongside the nutritional and immunological ones.

Some formula manufacturers have begun adding synthetic HMOs — typically 2'-fucosyllactose (2'-FL), the most abundant HMO — to their products. Early studies suggest that these additions can partially shift the formula-fed infant microbiome toward a more *Bifidobacterium*-enriched profile. However, the full complexity of the HMO repertoire — more than 200 distinct molecules, varying between individual mothers and across the course of lactation — is far beyond what any formula currently replicates.

Breastmilk also contains living bacteria — an estimated 600 to 800 species have been detected, though most are present at very low abundance. The dominant genera include *Staphylococcus*, *Streptococcus*, *Lactobacillus*, and *Bifidobacterium*. How these bacteria get into the milk remains incompletely understood. Some are thought to originate from the skin of the breast and areola. Others may be transported from the mother's gut through an *entero-mammary pathway* — a hypothesised route in which immune cells in the mother's intestinal wall engulf gut bacteria and carry

them through the lymphatic system to the mammary gland. This pathway has been demonstrated in animal models, and there is circumstantial evidence for it in humans, though it remains an active area of investigation.

The picture that emerges is one of breathtaking evolutionary engineering. The mother does not simply deliver her infant into a microbial world and hope for the best. She actively provisions her child with a starter culture of organisms, feeds those organisms with a bespoke food supply, and simultaneously provides antibodies and immune factors that protect the infant while its own immune system is still learning to function. The breastfed infant's microbiome is not acquired by accident. It is *curated*.

3.4 The First Thousand Days: Assembling a Stable Ecosystem

The period from conception to roughly two years of age — often called the *first thousand days* — is the most dynamic and consequential phase of microbiome development. During this window, the infant gut undergoes a transformation as dramatic as any ecological succession: from a near-sterile environment to a complex, diverse, functionally mature community of hundreds of species.

The process is not random, and it is not instantaneous. It unfolds in a broadly predictable sequence, though with substantial individual variation.

In the first days and weeks, the gut is colonised primarily by *facultative anaerobes* — organisms that can tolerate oxygen. These include *Escherichia coli*, *Enterococcus*, and *Staphylococcus* species. Their metabolic activity consumes the residual oxygen in the neonatal gut, progressively lowering the oxygen tension and creating the anaerobic environment that the next wave of colonisers requires. This is ecological succession in miniature: the pioneer species literally change the environment, making it habitable for the organisms that will eventually replace them.

As the gut becomes increasingly anaerobic, *obligate anaerobes* begin to establish themselves. In breastfed infants, this niche is dominated by *Bifidobacterium*, as we have seen. In formula-fed infants, the anaerobic community is more heterogeneous. In either case, the transition from a facultative-anaerobe-dominated community to one dominated by obligate anaerobes is a universal feature of early gut colonisation.

The introduction of solid food — typically between four and six months of age — triggers the next major transition. Suddenly, the gut is presented with an entirely new array of substrates: plant polysaccharides, complex starches, proteins, and fats that were absent from a diet of milk alone. The microbial community responds accordingly. Species capable of fermenting complex plant carbohydrates — notably members of the *Bacteroides* and various *Clostridium* clusters — expand. *Bifidobacterium* abundance, no longer sustained by a diet rich in HMOs, begins to decline. The diversity of the community increases markedly. By the time a child is eating a varied diet of solid foods, the gut microbiome is beginning to resemble, in broad outline, the community it will carry into adulthood.

By approximately three years of age — though estimates range from two to five years depending on the study and the population — the gut microbiome reaches a state of relative maturity. Its composition, while still influenced by diet, illness, and environment, is recognisably adult-like in its overall structure. The major phyla are established. The core metabolic functions — fibre fermentation, vitamin synthesis, bile acid transformation, colonisation resistance — are operational. The immune

system, trained over three years of continuous microbial dialogue, has developed the sophisticated ability to tolerate resident organisms while remaining vigilant against genuine threats.

This window of development is consequential precisely because it is a window — a period of heightened sensitivity during which disruptions can have outsized effects. The developing microbiome is like wet concrete: still malleable, still being shaped. Once it sets, it is far more resistant to change. But while it is setting, perturbations — particularly severe ones — can alter the final form.

The most well-studied perturbation is antibiotic exposure. Multiple large cohort studies have shown that antibiotic use in the first two years of life is associated with altered microbiome composition that can persist for months or years. A single course of broad-spectrum antibiotics can reduce microbial diversity, eliminate sensitive species, and allow antibiotic-resistant organisms to expand into the vacated niches. In some cases, species lost to antibiotic treatment in infancy are never fully recovered.

The epidemiological data are broadly consistent with this concern. Early-life antibiotic exposure has been associated — in observational studies — with increased rates of childhood obesity, asthma, allergic disease, and inflammatory bowel disease. The associations are modest in size and, as always with observational data, cannot definitively establish causation. But they are consistent across studies, biologically plausible given what we know about the microbiome's role in immune and metabolic development, and worrying enough that paediatric guidelines increasingly emphasise antibiotic stewardship in young children.

This does not mean that antibiotics should never be given to infants. A child with bacterial meningitis or sepsis needs antibiotics, and the benefit of treatment in such cases vastly outweighs any microbiome concern. But it does mean that the routine prescription of antibiotics for minor, self-limiting infections in early childhood — viral upper respiratory infections, uncomplicated ear infections — deserves more careful scrutiny than it has traditionally received.

3.5 The Adult Microbiome: Stable but Not Static

By early adulthood, the gut microbiome has achieved a degree of stability that is, in some ways, remarkable. Longitudinal studies that have tracked individuals over months and years have shown that the overall composition of an adult's gut microbiome is surprisingly consistent — more similar to itself across time than it is to the microbiome of any other individual. Your gut microbiome is, in a meaningful sense, a fingerprint: unique to you, and recognisably yours over extended periods.

This stability is maintained by several reinforcing mechanisms. The established community of organisms occupies the available ecological niches, making it difficult for newcomers to gain a foothold — the colonisation resistance we discussed in Chapter 1 and will explore further in Chapter 5. The immune system, now fully trained, actively maintains the composition of the microbial community through a continuous dialogue of molecular signals — a topic we will examine in detail in Chapter 12. And the host's diet and lifestyle, which tend to be relatively consistent in adulthood, provide a stable selective environment.

But stable does not mean static. The adult microbiome fluctuates continuously in response to short-term perturbations. A course of antibiotics disrupts it. A change in diet shifts it. Travel to a different country exposes it to new organisms. Illness, stress, sleep disruption, and even the seasons can alter its composition on timescales of days to weeks.

What is striking is the resilience — the tendency of the microbiome to return to something resembling its original state after a perturbation ends. Stop the antibiotics, and the community begins to recover. Return to your usual diet after a holiday, and the microbial composition drifts back toward its baseline. This resilience is itself an ecological property: it reflects the diversity and functional redundancy of the community, and the strength of the interactions — competitive, cooperative, and antagonistic — that hold the ecosystem together.

However, resilience has limits. Severe or repeated perturbations can push the microbiome past a tipping point from which it does not fully recover. The concept of *alternative stable states* — borrowed from ecology, where it describes lakes that can flip from clear to turbid and remain turbid even after the original stressor is removed — has been applied to the gut microbiome. Some researchers believe that the loss of key species through repeated antibiotic exposure, sustained dietary impoverishment, or chronic illness can shift the microbiome into a different configuration — less diverse, less functionally capable, more susceptible to pathogen invasion — from which it does not spontaneously return. Whether and how often this occurs in practice remains one of the most important open questions in microbiome science.

3.6 Ageing and the Microbiome: The Final Transition

The microbiome does not remain static as we age. In the later decades of life, a gradual but measurable shift occurs — one that may have profound implications for the health of older adults.

The most consistent finding across studies of the elderly microbiome is a decline in diversity. The rich, complex community of adulthood becomes, on average, less species-rich in old age. The relative abundance of *Bifidobacterium* — the genus that dominated the infant gut — declines further. The balance between the major phyla shifts: some studies report increases in *Proteobacteria* (a phylum that includes many opportunistic pathogens) and decreases in some of the fibre-fermenting *Clostridium* clusters that are important producers of the anti-inflammatory short-chain fatty acid butyrate.

The causes of these changes are multiple and difficult to disentangle. Diet plays a role: older adults often eat less fibre, fewer varied foods, and more processed meals, particularly those in residential care. Medication is another factor: polypharmacy — the use of multiple medications simultaneously — is common in the elderly, and many non-antibiotic drugs have significant effects on the microbiome, as we will explore in Chapter 22. The ageing immune system (*immunosenescence*) may also contribute: as immune surveillance weakens, the finely tuned dialogue between host and microbiome may become less precise, allowing shifts in community composition that a younger immune system would have corrected.

The institutional environment matters, too. Studies comparing community-dwelling elderly with those in long-term residential care have found striking differences in microbiome composition. Residents of care homes tend to have lower microbial diversity and a microbiome profile associated with increased inflammation. It is difficult to determine the direction of causality — does the institutional environment cause the microbiome changes, or do the health conditions that lead to institutionalisation cause both? — but the correlation is robust.

What is most intriguing, and most concerning, is the association between microbiome changes in the elderly and the constellation of problems that define geriatric frailty: chronic low-grade inflammation (sometimes called *inflammageing*), sarcopenia (loss of muscle mass), cognitive decline, and

increased susceptibility to infection. The causal relationships here are, to be frank, not yet established. But the correlations are strong enough that the microbiome is increasingly viewed as both a potential biomarker of healthy ageing and a potential target for intervention — a topic we will return to in the later chapters of this book.

There is also a more optimistic perspective. If the age-related decline in microbiome diversity is driven largely by modifiable factors — diet, medication, physical activity, social environment — then it may, at least in part, be preventable or reversible. A growing body of research suggests that dietary interventions rich in fibre and fermented foods can partially restore microbial diversity in older adults. Physical exercise, even at moderate levels, is associated with a more diverse and functionally capable gut microbiome. And the emerging field of microbiome-targeted therapeutics may eventually offer more precise interventions for age-related dysbiosis.

The trajectory of the microbiome across a human lifetime, then, follows a broad arc: rapid colonisation and assembly in infancy, maturation and stabilisation in childhood, relative stability through adulthood, and gradual decline in old age. At every stage, the microbiome is shaped by the interplay between the organisms it contains, the immune system of its host, and the environment in which both exist. It is not a fixed endowment. It is a living, changing ecosystem — one that begins as an inheritance from your mother and ends as a reflection of your entire life.

3.7 Your Microbiome Is Not Just Yours

Before we leave this chapter, there is one more idea worth considering — one that connects the personal story told here to the evolutionary story told in Chapter 2.

When we speak of genetic inheritance, we usually mean the DNA you received from your parents: half from your mother, half from your father, combined at conception into a genome that is uniquely yours. This inheritance is vertical — it flows from parent to child — and it is fixed at the moment of fertilisation.

Your microbial inheritance is also vertical, but it is far more complex. You received your founding microbial community primarily from your mother — through birth, through breastfeeding, through skin contact, through the shared environment of early life. Your father, your siblings, your household, your pets, and eventually your wider social world all contributed organisms to your developing microbiome. This inheritance was not fixed at a single moment. It accumulated over months and years, shaped by the ecological dynamics of succession, competition, and co-operation among the organisms themselves.

And here is the part that should give us pause. Your microbial inheritance is also, in some sense, *cultural*. The microbiome you carry reflects not just your mother's biology, but her diet, her environment, her medical history, and the microbial environment of the community in which you were born. Populations that eat traditional diets rich in plant fibre carry different microbiomes from populations that eat industrialised Western diets. These differences are transmitted vertically — from mother to child — and maintained horizontally — through shared food, shared environment, shared ways of life.

This means that when a population's diet changes — when traditional foods are replaced by processed ones, when dietary fibre disappears, when antibiotic use becomes widespread — the microbial inheritance itself can change. Not in a single generation, but across generations, as mothers

pass on diminished microbial communities to their children, who pass on further diminished communities to theirs. The microbiologist Martin Blaser has called this the *disappearing microbiota hypothesis* — the idea that the microbiome of industrialised populations has been progressively depleted over generations by the combined pressures of antibiotics, dietary change, sanitation, and caesarean delivery, and that some of the species lost may never be recovered.

This is a sobering thought: that we may be losing, with each generation, not just cultural knowledge or biodiversity, but the microbial partners that helped build us and that our bodies evolved to expect.

Whether this hypothesis is correct in its strongest form remains to be determined. But it underscores a theme that runs through this entire book: the microbiome is not a luxury. It is not an accessory. It is an integral part of what it means to be a human organism, and its loss — whether in an individual or across a population — has consequences we are only beginning to understand.

In the next chapter, we will look at the tools that have made all of this knowledge possible — the technologies that lifted the veil on the microbial world and revealed the ecosystem within.

Chapter 3 — Key References

[To be populated with full citations during manuscript preparation]

- Dominguez-Bello, M.G. et al. (2010). Delivery mode shapes the acquisition and structure of the initial microbiota across multiple body habitats in newborns. *Proceedings of the National Academy of Sciences*, 107(26), 11971–11975.
- Bäckhed, F. et al. (2015). Dynamics and stabilization of the human gut microbiome during the first year of life. *Cell Host & Microbe*, 17(5), 690–703.
- Stewart, C.J. et al. (2018). Temporal development of the gut microbiome in early childhood from the TEDDY study. *Nature*, 562(7728), 583–588.
- Bode, L. (2012). Human milk oligosaccharides: every baby needs a sugar mama. *Glycobiology*, 22(9), 1147–1162.
- Sela, D.A. & Mills, D.A. (2010). Nursing our microbiota: molecular linkages between bifidobacteria and milk oligosaccharides. *Trends in Microbiology*, 18(7), 298–307.
- Pannaraj, P.S. et al. (2017). Association between breast milk bacterial communities and establishment and development of the infant gut microbiome. *JAMA Pediatrics*, 171(7), 647–654.
- Yassour, M. et al. (2016). Natural history of the infant gut microbiome and impact of antibiotic treatment on bacterial strain diversity and stability. *Science Translational Medicine*, 8(343), 343ra81.
- Blaser, M.J. (2014). *Missing Microbes: How the Overuse of Antibiotics Is Fueling Our Modern Plagues*. Henry Holt and Company.
- Sonnenburg, E.D. et al. (2016). Diet-induced extinctions in the gut microbiota compound over generations. *Nature*, 529(7585), 212–215.
- O'Toole, P.W. & Jeffery, I.B. (2015). Gut microbiota and aging. *Science*, 350(6265), 1214–1215.
- Claesson, M.J. et al. (2012). Gut microbiota composition correlates with diet and health in the elderly. *Nature*, 488(7410), 178–184.
- de Goffau, M.C. et al. (2019). Human placenta has no microbiome but can contain potential pathogens. *Nature*, 572(7769), 329–334.
- Cunnington, A.J. et al. (2016). “Vaginal seeding” of infants born by caesarean section. *BMJ*,

352, i227.

- Korpela, K. et al. (2018). Intestinal microbiome is related to lifetime antibiotic use in Finnish pre-school children. *Nature Communications*, 9, 1517.
- Shao, Y. et al. (2019). Stunted microbiota and opportunistic pathogen colonization in caesarean-section birth. *Nature*, 574(7776), 117–121.

Chapter 4: Tools of Discovery — How We Study What We Can't See

The previous three chapters told a story of partnership — between humans and microbes, between viruses and genomes, between mothers and the microbial communities they pass to their children. But telling that story required knowing things that, for most of human history, no one could possibly have known. How do you study organisms too small to see? How do you identify thousands of species living in a single gram of intestinal content? How do you determine what those species are doing — not in a laboratory flask, but inside a living human body?

The history of microbiome science is, in large part, a history of tools. Every major advance in our understanding of the microbial world has followed an advance in our ability to observe it. Van Leeuwenhoek's hand-ground lenses revealed the existence of bacteria. Koch's agar plates and staining techniques allowed individual species to be isolated and identified. And the molecular revolution — beginning in the 1970s and accelerating dramatically in the 2000s — cracked open a world that culture-based methods had barely hinted at.

This chapter is a guide to that toolkit. It is, necessarily, more technical than the chapters that came before. But the technologies described here are not abstractions. They are the instruments through which everything else in this book was discovered. Without them, the microbiome would still be invisible.

4.1 From Petri Dishes to DNA: The Culture Problem

For over a century after Koch, microbiology rested on a single foundational technique: culturing. To identify a microorganism, you grew it. You took a sample — from soil, from water, from a patient's blood or stool — and you spread it on a nutrient medium, incubated it under controlled conditions, and waited. If a species could grow on the medium you provided, under the atmospheric conditions you chose, at the temperature you set, within the timeframe you allowed, it would form a visible colony. You could then isolate that colony, stain it, examine it under a microscope, test its biochemical reactions, and assign it a name.

This approach was spectacularly successful for identifying pathogens. The agents of tuberculosis, cholera, diphtheria, typhoid, plague, and dozens of other diseases were isolated and characterised using culture methods. Clinical microbiology — the diagnostic discipline that tells a physician which bacterium is causing a patient's pneumonia or urinary tract infection — still relies heavily on culture today, and for good reason: it works, and it provides a living isolate that can be tested for antibiotic susceptibility.

But for studying microbial communities — the teeming, diverse, interdependent ecosystems that exist in environments like the human gut — culture was a disaster.

The problem was recognised as early as 1985, when the microbiologists James Staley and Allan Konopka gave it a name: the *great plate count anomaly* [REF:staley1985]. When you examined

a sample of natural water or soil under a microscope, you could count the cells directly. When you then tried to grow those same cells on standard laboratory media, only a tiny fraction formed colonies — typically 0.1 to 1 per cent of the total. The vast majority of the organisms you could see were organisms you could not *grow*.

The anomaly turned out to be even more extreme in the human gut. The large intestine is a deeply anaerobic environment — almost entirely devoid of oxygen. Many of the dominant gut bacteria are *obligate anaerobes*: organisms that not only do not require oxygen but are actively killed by it. Exposing a stool sample to air, even briefly, can destroy the very species you are trying to study. Others require specific nutrients, particular pH ranges, the metabolic products of neighbouring organisms, or growth conditions so precise that no standard laboratory protocol could reproduce them.

The result was a massive blind spot. For decades, microbiologists knew that the gut contained bacteria — they could see them, and they could culture some of them, including clinically important species like *Escherichia coli*, *Bacteroides fragilis*, and *Clostridium perfringens*. But the full diversity of the community was invisible. Estimates now suggest that traditional culture methods could recover perhaps 20 to 30 per cent of the species present in the human gut — and far less in environments like soil or seawater.

This meant that the majority of microbial life on Earth — including the majority of microbial life inside your body — was unknown. Not merely uncharacterised. *Unknown*. Entire phyla of bacteria existed that had never been grown in a laboratory, never been given a name, never been studied in any detail. They were the dark matter of biology.

The solution would come not from better culture methods, but from an entirely different approach: reading DNA directly from environmental samples, without the need to grow anything at all.

4.2 The Molecular Rosetta Stone: 16S Ribosomal RNA

The conceptual breakthrough that made culture-independent microbiology possible came from an unlikely source: a quest to understand the deepest branches of the tree of life.

In 1977, Carl Woese and George Fox, working at the University of Illinois, published a paper in the *Proceedings of the National Academy of Sciences* that would reshape biology [REF:woese1977]. They had been comparing the sequences of a molecule called 16S ribosomal RNA (16S rRNA) — a component of the ribosome, the cellular machinery that translates genetic information into protein — across a wide range of organisms. The ribosome is universal. Every living cell on Earth possesses one. And because the ribosome's function is so fundamental, its structure has been conserved across billions of years of evolution. The 16S rRNA gene changes slowly enough that it retains recognisable similarities even between very distantly related organisms, yet fast enough that it accumulates differences between species that can be used to tell them apart.

Woese and Fox's comparison revealed something no one had expected: what had been classified as bacteria were actually two fundamentally distinct domains of life. One domain — the *Bacteria* — included all of the familiar species that microbiologists had been studying for a century. The other — the *Archaea* — looked superficially similar under a microscope but were as genetically distant from bacteria as either group was from plants and animals. The discovery of the Archaea as a separate domain of life, based entirely on molecular sequence comparison, was one of the great conceptual

revolutions in twentieth-century biology.

But the 16S rRNA gene had another property that would prove even more consequential for microbiome science. Because it is present in all bacteria and archaea, it functions as a universal identifier — a molecular barcode. The approximately 1,500 nucleotides of the 16S rRNA gene contain regions that are highly conserved (nearly identical across all bacteria) interspersed with regions that are highly variable (different between species, and sometimes between strains within a species). The conserved regions provide convenient anchoring points where universal PCR primers can bind. The variable regions provide the discriminating information that allows you to identify what you have caught.

Think of it like a passport. Every passport has the same basic format — the same layout, the same fields — which is what allows border agents to read any passport regardless of nationality. But the information in those fields — the name, the photograph, the passport number — is what distinguishes one traveller from another. The 16S rRNA gene works the same way. The conserved regions are the format; the variable regions are the identity.

This insight led to a simple but revolutionary idea, articulated most clearly by the microbial ecologist Norman Pace in the 1990s [REF:pace1997]. If you could extract DNA directly from an environmental sample — a scoop of soil, a millilitre of seawater, a gram of faeces — and amplify just the 16S rRNA genes using those universal primers, you could identify the bacteria present without ever growing them. You would not need a Petri dish or an incubator. You would not need to know what medium each species required, or what temperature it preferred, or whether it could tolerate oxygen. You would simply read its molecular barcode.

The technique is called *16S rRNA gene sequencing*, sometimes referred to as *amplicon sequencing*, and it arrived just as a second revolution was getting underway: the development of high-throughput DNA sequencing technologies that could process thousands or millions of sequences in a single run.

4.3 The First Census: 16S Sequencing in Practice

The basic workflow of a 16S rRNA gene survey, stripped to its essentials, runs as follows.

First, you collect your sample. In human microbiome studies, this is most commonly a stool sample (as a proxy for the gut microbiota), though swabs of the skin, oral cavity, vaginal tract, or other body sites are also routine. The sample must be collected, stored, and transported under conditions that preserve the DNA while minimising contamination — a non-trivial logistical challenge, particularly for large cohort studies involving hundreds or thousands of participants.

Second, you extract the DNA. This involves breaking open the microbial cells — using chemical, enzymatic, or mechanical means — and purifying the released DNA from everything else in the sample: proteins, lipids, bile salts, dietary residues, and human cells. The method used to lyse (break open) cells matters enormously. Some bacteria, particularly Gram-positive species with thick cell walls, are much harder to crack open than others. If your lysis method is too gentle, you will underrepresent the species with tough cell walls; if it is too harsh, you may shear the DNA into fragments too small to sequence. Different DNA extraction kits can produce significantly different community profiles from the same starting material — a source of variability that must be carefully controlled in any comparative study [REF:jones2015].

Third, you amplify the 16S rRNA gene using PCR (polymerase chain reaction). The primers are designed to bind to conserved regions flanking one or more of the variable regions — most commonly the V3–V4 or V4 region alone. The amplified DNA fragments — one for each bacterium whose 16S gene was captured — are then prepared for sequencing.

Fourth, you sequence. In the early days of microbiome research, this meant Sanger sequencing — the gold standard method developed by Frederick Sanger in the 1970s, capable of reading individual DNA fragments of up to 1,000 base pairs with high accuracy, but expensive and slow. A typical early study might sequence a few hundred clones from a sample, enough to identify the dominant species but far too few to capture the full diversity of a complex community.

The landscape changed dramatically with the arrival of *next-generation sequencing* (NGS) platforms in the mid-2000s [REF:wensel2022]. Technologies developed by 454 Life Sciences (later Roche), Illumina, and Ion Torrent could produce hundreds of thousands to millions of short sequence reads from a single run. The Illumina platform, which came to dominate the field, uses a chemistry called sequencing-by-synthesis: millions of DNA fragments are attached to a glass surface, amplified into clusters, and then read simultaneously by detecting the incorporation of fluorescently labelled nucleotides one base at a time. The result is an avalanche of data — millions of short reads, each 150 to 300 base pairs long — at a fraction of the cost of Sanger sequencing.

Finally, you analyse the data. Each sequence read is compared against reference databases — curated collections of known 16S rRNA gene sequences, such as the SILVA database, the Ribosomal Database Project (RDP), or the Greengenes database — to determine which organism it came from. Sequences that are sufficiently similar to each other (traditionally 97 per cent identical, though more recent methods use exact sequence variants, as we will discuss) are grouped together into *operational taxonomic units* (OTUs) — a practical proxy for species when formal taxonomic classification is not possible.

The output of a 16S rRNA gene survey is a table: a matrix showing which OTUs (or exact sequence variants) are present in each sample, and at what relative abundance. This table is the raw material from which all downstream analyses flow. You can calculate *alpha diversity* — how many different species are present in a single sample. You can calculate *beta diversity* — how similar or different two samples are in their community composition. You can test whether a particular species is more or less abundant in patients with a disease compared to healthy controls. You can track changes in community composition over time, across body sites, or in response to interventions.

The Human Microbiome Project, which we introduced in Chapter 1, was the first large-scale application of this approach to the human body. Using 16S rRNA gene sequencing alongside shotgun metagenomics, the project characterised the microbial communities at 18 body sites in 242 healthy adults, producing the first comprehensive map of the healthy human microbiome [REF:hmp2012].

But 16S sequencing, powerful as it was, had important limitations — limitations that would drive the development of the next generation of tools.

4.4 Beyond the Barcode: Shotgun Metagenomics

The 16S rRNA gene is an identifier, not a description. It tells you *who* is present in a community, but it tells you almost nothing about *what they are doing*. Knowing that your gut contains a species of *Bacteroides* is useful, but it does not tell you which metabolic pathways that species is running,

which genes it is expressing, whether it carries antibiotic resistance genes, or how it interacts with its neighbours. For that, you need to read more than just the barcode. You need to read the whole genome — or rather, all of the genomes.

This is the approach known as *shotgun metagenomics*, and it works on a fundamentally different principle from 16S sequencing [REF:wensel2022].

Instead of amplifying a single target gene with specific primers, shotgun metagenomics takes all of the DNA in a sample — bacterial, archaeal, fungal, viral, and human — and fragments it randomly into short pieces. These fragments are then sequenced, producing millions of reads that are collectively a jumbled snapshot of every genome present in the sample. The computational challenge is to reassemble these fragments, or at least to assign them to known organisms and known gene functions, using reference databases and sophisticated bioinformatic algorithms.

The analogy is often drawn to a shredder. Imagine taking the complete works of Shakespeare, Tolstoy, a Cantonese cookbook, and a plumbing manual, shredding them all together into strips of five to ten words, and then trying to reconstruct the original texts — or at least to figure out which languages were represented and what topics were covered. Shotgun metagenomics is doing essentially this, except that the “texts” are genomes, the “strips” are sequence reads, and the reconstruction is performed by algorithms that can compare each strip against a vast library of known texts.

The payoffs are substantial. Shotgun metagenomics provides not just taxonomic identification but functional information: which metabolic genes are present, which antibiotic resistance determinants are circulating in the community (the *resistome*), which virulence factors are encoded. It can identify organisms at the species or even the strain level — a resolution that 16S sequencing, which typically operates at the genus level, cannot match. And it can detect organisms that 16S sequencing misses entirely: viruses, which have no ribosomal RNA gene; some fungi and parasites, whose ribosomal genes may not amplify well with standard bacterial primers; and even novel organisms with no close relatives in existing databases.

The landmark MetaHIT study (Metagenomics of the Human Intestinal Tract), published in *Nature* in 2010 by Junjie Qin and colleagues, demonstrated the power of this approach [REF:qin2010]. By performing shotgun metagenomic sequencing of faecal samples from 124 European individuals, the team assembled a catalogue of 3.3 million non-redundant microbial genes — roughly 150 times more genetic information than the human genome contains. They estimated that the cohort collectively harboured between 1,000 and 1,150 prevalent bacterial species, with each individual carrying at least 160 such species.

Shotgun metagenomics has its own limitations, however. It is more expensive than 16S sequencing, requires significantly more computational resources to analyse, and can be confounded by host DNA — in a stool sample, perhaps 5 to 10 per cent of the total DNA may be human; in a nasal swab or biopsy, the proportion can be 90 per cent or more, requiring either depletion of host DNA before sequencing or vastly more sequencing depth to capture the microbial fraction.

There is also a subtler problem. Shotgun metagenomics tells you which genes are *present* in a community, but not which genes are *active*. A bacterium carrying a gene for antibiotic resistance is not the same as a bacterium actively expressing that gene. The genome is a parts list, not an operating manual. To understand what the microbiome is actually doing at any given moment, you need to look not at the DNA but at its products.

4.5 What Are They Actually Doing? Metatranscriptomics, Metabolomics, and Proteomics

The DNA in a microbial community is relatively stable. The genes present today will, for the most part, still be present tomorrow. But gene *expression* — which genes are being actively transcribed into messenger RNA, and which proteins and metabolites are being produced — changes rapidly in response to diet, medication, host immune signals, and the activities of neighbouring organisms. To capture this dynamic layer, researchers have developed a suite of complementary approaches, collectively known as the *multi-omics* toolkit.

Metatranscriptomics extracts and sequences the RNA, rather than the DNA, from a microbial community. Because messenger RNA (mRNA) is produced only from genes that are being actively transcribed, the metatranscriptome provides a snapshot of gene expression — a readout of what the community is doing right now, rather than what it has the potential to do. Metatranscriptomics has revealed that gene expression patterns in the gut microbiome shift dramatically in response to dietary changes, illness, and circadian rhythms, even when the underlying DNA composition of the community remains stable. The technical challenges are considerable: RNA is far less stable than DNA, degrades rapidly after sample collection, and requires careful handling and specialised extraction protocols to preserve.

Metabolomics takes a different approach entirely. Rather than sequencing nucleic acids, it uses analytical chemistry — typically mass spectrometry or nuclear magnetic resonance (NMR) spectroscopy — to identify and quantify the small molecules (metabolites) present in a sample. The metabolome is, in a sense, the final output of all the genomic activity: the short-chain fatty acids, bile acid derivatives, amino acid metabolites, vitamins, and signalling molecules that the microbiome actually produces and that interact with the host. When we say that the gut microbiome “communicates with the brain” or “regulates the immune system,” it is largely through metabolites that this communication occurs.

Metabolomics has revealed that the microbiome produces thousands of small molecules, many of which were previously unknown and whose biological functions are still being characterised. Some of these are familiar: butyrate, propionate, and acetate — the short-chain fatty acids produced by bacterial fermentation of dietary fibre — are among the most studied metabolites in microbiome science. Others are less familiar but potentially just as important: trimethylamine (TMA), produced by gut bacteria from dietary choline and carnitine and converted by the liver to trimethylamine N-oxide (TMAO), a metabolite implicated in cardiovascular disease; or indole derivatives, produced from the amino acid tryptophan, which modulate intestinal barrier function and immune responses.

Metaproteomics — the large-scale identification of proteins in a microbial community — completes the picture by showing which proteins are actually being manufactured. While technically demanding and still less widely used than the other approaches, metaproteomics can reveal which enzymes are actively at work in the gut, providing information that neither genomics nor transcriptomics alone can supply (a gene may be transcribed into mRNA without the mRNA ever being translated into functional protein).

The most powerful studies now combine multiple omics layers — metagenomics and metabolomics, for example, or metagenomics and metatranscriptomics — to build integrated pictures of microbial community function. This *multi-omics* approach, while computationally challenging, is beginning to reveal how microbial communities operate as systems rather than as collections of individual species [REF:knight2018].

4.6 ITS Sequencing and the Mycobiome Blind Spot

There is a conspicuous absence in everything described so far: fungi.

The 16S rRNA gene is specific to bacteria and archaea. Fungi have ribosomes, of course, but their ribosomal RNA genes are structurally different and are not amplified by the universal bacterial primers used in standard 16S surveys. For decades, this meant that most microbiome studies simply did not detect fungi. The organisms were there — *Candida* in the gut, *Malassezia* on the skin, *Aspergillus* in the lungs — but the standard tools could not see them. The mycobiome was, in effect, methodologically invisible.

The fungal equivalent of the 16S gene is the *internal transcribed spacer* (ITS) region — a stretch of DNA that lies between the ribosomal RNA genes in the fungal genome. Like the bacterial 16S gene, the ITS region is present in all fungi, varies enough between species to be useful for identification, and can be amplified with broadly specific primers. ITS sequencing has become the standard method for characterising fungal communities and is the primary tool used in mycobiome research.

But ITS sequencing brings its own challenges. The ITS region is more variable in length than the 16S gene, which can complicate sequencing and bioinformatic analysis. Reference databases for fungal ITS sequences are less complete than their bacterial counterparts, meaning that a larger proportion of sequences cannot be confidently assigned to known species. And the sheer dominance of bacteria in most body sites means that, even in samples where fungi are present, they represent a tiny fraction of the total microbial biomass — often less than 0.1 per cent — making them difficult to detect without targeted amplification.

More recently, long-read sequencing technologies — particularly those developed by Oxford Nanopore Technologies and Pacific Biosciences — have offered new approaches to fungal identification that can read through the entire ITS region (and beyond) in a single pass, avoiding some of the biases inherent in short-read sequencing of variable-length amplicons [REF:dandrea2021].

The belated recognition that fungi are a consistent, functionally important component of the human microbiome — interacting with bacteria, modulating the immune system, and expanding opportunistically when bacterial competitors are removed by antibiotics — is one of the reasons we have devoted an entire chapter to the mycobiome later in this book. The tools shaped what we could see, and for years, fungi were outside the frame.

4.7 Gnotobiotic Animals: What Happens When You Remove All the Microbes?

Sequencing technologies can tell you who is present and what genes they carry. But they cannot, on their own, tell you what those organisms *do to the host*. For that, you need experiments — and the most powerful experimental model in microbiome science is the *gnotobiotic* animal.

The term *gnotobiotic* comes from the Greek *gnotos* (known) and *bios* (life): it refers to an organism whose microbial status is entirely known and controlled. In practice, this usually means *germ-free* animals — mice, rats, or other species that have been born and raised in completely sterile environments, with no microbial colonisation of any kind. These animals are delivered by caesarean section into sterile isolators, fed autoclaved food and water, and maintained in sealed chambers with filtered air for their entire lives.

Germ-free mice are strange creatures. They have underdeveloped immune systems — their gut-associated lymphoid tissue is atrophied, their T cell populations are abnormal, and their ability to mount immune responses to pathogens is severely compromised. Their intestinal walls are thinner. Their caecum — the pouch at the junction of the small and large intestines — is massively enlarged, bloated with undigested mucus that the absent microbiome would normally have metabolised. They consume more food but weigh less than their conventionally colonised counterparts (a finding that complicated early attempts to study the microbiome's role in obesity). Their behaviour is altered: germ-free mice tend to show reduced anxiety-like behaviours in standard behavioural tests, an observation that helped catalyse the study of the gut-brain axis.

The experimental power of gnotobiotic systems lies in what you can do next. You can take a germ-free mouse and introduce a single known bacterial species — creating a *monocolonised* animal — and study what that one species does to the host in isolation, free from the confounding effects of every other organism. You can introduce a defined community of two, five, or twenty species — creating what is sometimes called a *synthetic community* or *SynCom* — and study their interactions. Or, most dramatically, you can take a faecal sample from a human subject and transplant it into a germ-free mouse, creating what is called a *human microbiota-associated* (HMA) mouse — an animal carrying a humanised microbiome.

HMA mice have been used to address some of the biggest questions in microbiome science. The landmark experiments by Jeffrey Gordon's laboratory at Washington University, beginning in 2006, showed that transferring the gut microbiota from obese mice into germ-free recipients caused the recipients to gain more fat than those colonised with microbiota from lean donors — even when both groups ate the same diet [REF:ley2006]. Later experiments transferred human microbiota from obese and lean twin pairs into germ-free mice, replicating the phenotype. These studies were instrumental in establishing the idea that the microbiome itself can influence body weight and metabolic function.

But gnotobiotic experiments also have limitations that are important to acknowledge — and that have led to significant criticism of how their results are sometimes interpreted.

A 2020 systematic review by Jens Walter, Anissa Armet, Brett Finlay, and Fergus Shanahan, published in *Cell*, examined all published studies using HMA mice and found a striking result: 95 per cent of studies reported that transferring a disease-associated human microbiome into germ-free mice reproduced the disease phenotype in the animals [REF:walter2020]. The authors argued that this success rate is implausibly high and likely reflects publication bias (studies that fail to show a phenotype are less likely to be published), methodological issues (germ-free mice are so physiologically abnormal that almost any colonisation event produces measurable changes), and overinterpretation (correlation between a human disease microbiome and a mouse phenotype does not prove that the microbiome caused the human disease).

The critique is a valuable corrective. Gnotobiotic models remain indispensable — they are the closest thing we have to a controlled experiment on the microbiome's effects on a mammalian host. But their results must be interpreted cautiously, and the leap from “this microbiome produces this phenotype in a germ-free mouse” to “this microbiome causes this disease in humans” is larger than many published papers have acknowledged.

4.8 Computational Challenges: Bioinformatics, Big Data, and the Role of AI

Every technology described in this chapter produces data — and the quantities are staggering.

A single shotgun metagenomic sequencing run on an Illumina NovaSeq platform can generate several billion reads. A large cohort study — sequencing hundreds or thousands of samples, each producing millions of reads — generates terabytes of raw data that must be stored, processed, quality-filtered, assembled, annotated, and statistically analysed. The computational infrastructure required to handle this data is on a scale more familiar to astrophysics or particle physics than to traditional biology.

The bioinformatic pipeline — the sequence of computational steps that converts raw sequence reads into biological insight — is itself a complex and evolving discipline. For 16S sequencing, the pipeline involves quality filtering (removing low-quality reads and sequencing artefacts), chimera detection (identifying and removing artificial sequences created when two different template molecules are joined during PCR), taxonomic classification (assigning each read to a taxon using reference databases), and statistical analysis. For shotgun metagenomics, the pipeline adds read assembly (piecing together overlapping short reads into longer contiguous sequences, or *contigs*), gene prediction (identifying open reading frames within the assembled contigs), functional annotation (assigning predicted genes to known functional categories), and *binning* (grouping contigs that likely come from the same organism into *metagenome-assembled genomes*, or MAGs).

Each of these steps involves choices — which algorithm to use, which parameters to set, which databases to query — and those choices can significantly influence the results. Two different bioinformatic pipelines applied to the same raw data can produce different taxonomic profiles, different diversity estimates, and different statistical conclusions. This is not a failure of science; it is a reflection of the genuine difficulty of inferring biological reality from fragmentary sequence data. But it does mean that methodological transparency — reporting exactly which tools and parameters were used, and ideally making raw data publicly available for reanalysis — is essential.

One important methodological advance deserves specific mention. For years, the standard practice in 16S sequencing was to cluster similar sequences into operational taxonomic units at a 97 per cent identity threshold — a pragmatic approach that grouped sequences into approximate “species” but discarded the fine-scale variation that might distinguish closely related strains. Beginning around 2016, a new class of algorithms — including DADA2, Deblur, and UNOISE — replaced OTU clustering with *amplicon sequence variants* (ASVs), also called *exact sequence variants* (ESVs). These algorithms use statistical error models to correct for sequencing errors and resolve the data to individual, exact sequences — distinguishing variants that differ by as little as a single nucleotide [REF:knight2018]. The shift from OTUs to ASVs has improved reproducibility (the same biological sequence produces the same ASV regardless of which study or which sequencing run it came from), increased resolution, and enabled more precise ecological analyses.

Machine learning and artificial intelligence are increasingly being applied to microbiome data. Neural networks, random forests, and other algorithms can detect patterns in complex microbial community data that are invisible to conventional statistical methods — identifying combinations of species whose abundances predict disease status, or discovering non-obvious correlations between dietary patterns and microbiome composition. These approaches are particularly valuable when dealing with the high dimensionality of microbiome datasets: a typical study might measure the abundances of hundreds of microbial species across dozens of clinical and dietary variables, creating a matrix too complex for traditional hypothesis-driven analysis.

However, machine learning applied to microbiome data is subject to the same pitfalls that affect its application in other fields: overfitting (finding patterns in the training data that do not generalise to new data), confounding (learning associations that reflect underlying variables — age, diet, medication use — rather than the biological relationship of interest), and lack of interpretability (producing accurate predictions without explaining *why* they are accurate). The most rigorous studies use cross-validation, independent replication cohorts, and careful attention to potential confounders — practices that are becoming standard but are not yet universal.

4.9 The Sequencing Revolution: From Months to Minutes

The technologies described so far have followed a striking trajectory: faster, cheaper, more accessible. It is worth pausing to appreciate the scale of this change, because it is transforming not just research but the potential for clinical application.

The Human Genome Project, completed in 2003, sequenced a single human genome at a cost of approximately \$2.7 billion over thirteen years. By 2025, a human genome could be sequenced in hours for less than \$200. The reduction in sequencing cost has outpaced even Moore's Law — the observed doubling of computing power every two years that has driven the technology industry for half a century. Sequencing costs have fallen by roughly a factor of ten million since the early 2000s.

For microbiome research, this has been transformative. Studies that would have been prohibitively expensive a decade ago — sequencing the gut microbiome of 10,000 individuals, for example, or performing longitudinal sequencing of a single individual daily for a year — are now feasible. Large population-based cohorts, such as the American Gut Project, the Dutch LifeLines-DEEP cohort, and the UK Biobank (which has added microbiome sampling to its protocol), are generating microbiome datasets of unprecedented scale.

A more recent development — nanopore sequencing, developed by Oxford Nanopore Technologies — introduces a fundamentally different approach. Rather than reading DNA by detecting the incorporation of fluorescent nucleotides (as Illumina platforms do), nanopore sequencing threads single strands of DNA through a protein nanopore embedded in a membrane and detects the changes in electrical current as each nucleotide passes through the pore. The technology can read extremely long fragments of DNA — tens of thousands of base pairs in a single read, compared to the 150-to-300-base-pair reads of Illumina platforms — which is valuable for resolving complex genomic regions and for assembling complete microbial genomes from metagenomic data [REF:mitsuhashi2017].

Perhaps most remarkably, nanopore sequencing can be performed on a device the size of a USB flash drive — the MinION — which connects to a laptop and can generate results in real time. This portability opens possibilities that bench-top sequencers cannot: microbiome analysis in field settings, in remote clinics, and potentially at the point of care. While nanopore sequencing currently has a higher per-read error rate than Illumina, the gap is closing with each new chemistry iteration, and for many applications — species identification, resistance gene detection, rapid outbreak characterisation — the speed and portability outweigh the accuracy trade-off.

4.10 The Return of Culture: Culturomics

Given all the power of sequencing-based approaches, it might seem that traditional culture has nothing left to contribute. But a movement that began in Marseille, France, has demonstrated that culture is far from dead — it just needed to be reimaged.

The approach, developed by Jean-Christophe Lagier, Didier Raoult, and colleagues at the IHU Méditerranée Infection, is called *culturomics* [REF:lagier2018]. The concept is simple in principle, if laborious in execution: instead of growing bacteria on a handful of standard media under a single set of conditions, culturomics exposes each sample to dozens or even hundreds of different culture conditions — different media, different atmospheric compositions, different temperatures, different incubation times — in a systematic effort to coax into growth species that standard methods miss. Each colony that grows is then identified rapidly using *MALDI-TOF mass spectrometry* — a technique that identifies bacteria by their protein fingerprint in minutes, at a cost of a few cents per identification. Novel species that cannot be matched to existing databases are subjected to genome sequencing.

The results have been impressive. Since its introduction in 2012, culturomics has led to the isolation of hundreds of new bacterial species from the human gut, including many that had previously been known only from their DNA sequences. By combining culturomics with metagenomic sequencing, researchers have been able to put faces to names — to obtain living isolates of organisms that the molecular revolution had detected but could not grow, and to study their biology in the laboratory.

Culturomics also provides something that sequencing alone cannot: a living organism. You can sequence a genome from a metagenomic dataset, but you cannot perform biochemical assays, test antibiotic susceptibility, or study host-microbe interactions without a cultured isolate. For translational applications — developing probiotics, testing potential pathogens, understanding metabolic pathways — culture remains indispensable.

The revival of culture is a useful reminder that progress in science is not always linear. Technologies do not simply replace one another; they complement one another. The most complete picture of a microbial community comes from the integration of multiple approaches — molecular and culture-based, genomic and metabolic, computational and experimental.

4.11 Limitations and Pitfalls: How to Read Microbiome Science Critically

This chapter has described an impressive toolkit, and it would be easy to come away with the impression that microbiome science has the tools to answer any question. It does not. The limitations are real, and understanding them is essential for reading the rest of this book — and for evaluating the microbiome claims you will encounter in the media and the marketplace.

Contamination. All sequencing-based methods are vulnerable to contamination, and the problem is most acute in samples with low microbial biomass — blood, placental tissue, amniotic fluid, brain tissue, tumour biopsies. As the study by Salter and colleagues demonstrated in 2014, DNA extraction kits, PCR reagents, and laboratory water supplies can contain trace amounts of bacterial DNA from the manufacturing process [REF:salter2014]. In high-biomass samples like stool, this contamination is overwhelmed by the genuine microbial DNA in the sample and is negligible. In low-biomass samples, the contaminating DNA can constitute the majority of the sequences recovered, producing results that look like a microbial community but are entirely artefactual. The

sterile womb debate described in Chapter 3 was largely resolved by the recognition that many early reports of a placental microbiome were contamination artefacts.

The lesson for the reader is that any claim of a microbiome in a low-biomass body site — the brain, the blood, the placenta, the urinary tract — should be evaluated with particular caution, and with attention to whether the study included appropriate negative controls (samples processed identically but without biological material, to reveal the baseline contamination signature of the laboratory process).

Correlation versus causation. The most common finding in a microbiome study is an *association* — a statistical correlation between the presence or abundance of a particular organism and a clinical outcome. Patients with inflammatory bowel disease have different gut microbiomes from healthy controls. People with depression have altered microbial profiles. Obese individuals harbour different communities from lean individuals. These associations are real and reproducible. But they do not, in themselves, prove that the microbial differences caused the disease.

This distinction — between two things occurring together and one thing causing the other — sounds straightforward, but it trips up even sophisticated readers. So it is worth a brief detour into a famous example from outside medicine.

Across much of Western Europe over the second half of the twentieth century, populations of white storks (*Ciconia ciconia*) declined dramatically. Over the same period, birth rates fell. If you plot the two trends on the same graph — stork pairs on one axis, babies born on the other — the correlation is striking. In Germany, the coefficient was remarkably tight. A statistical test would confirm the association at high significance. And yet no reasonable person would conclude that the decline in storks *caused* the decline in human births, or that storks deliver babies. The explanation, of course, is that both trends were driven by a shared underlying cause: urbanisation. As rural landscapes were converted to cities and suburbs, stork nesting habitat disappeared *and* family sizes shrank as societies industrialised and women gained access to education and contraception. The storks and the babies had nothing to do with each other. They were both responding to the same world.

This is what statisticians call *confounding* — a hidden third variable that drives both of the things you are measuring, creating the illusion of a direct relationship between them. And confounding is rampant in microbiome research.

Consider a concrete example. Multiple studies have reported that patients with type 2 diabetes have altered gut microbiomes compared to healthy controls — different species, different proportions, different metabolic gene profiles. This is a genuine and reproducible finding. But patients with type 2 diabetes also differ from healthy controls in many other ways: they are more likely to be taking metformin (which has profound effects on the gut microbiome independent of diabetes), they tend to eat different diets, they exercise less, and they are more likely to be obese. Any of these factors — the medications, the diet, the activity levels, the excess weight — could be responsible for the observed microbial differences. The microbiome change might *cause* the diabetes. Or the diabetes might *cause* the microbiome change. Or — like the storks and the babies — both might be consequences of something else entirely, such as a high-sugar, low-fibre diet that simultaneously disrupts the microbiome and drives metabolic disease.

Disentangling these possibilities is one of the central challenges of the field, and it is genuinely difficult. There are three main relationships that a correlation between a microbiome change and a disease might reflect:

The microbiome change could be *causal* — the microbial shift actually drives or contributes to the

disease. This is the most exciting possibility and the one that dominates media headlines (“Gut Bacteria Cause Depression!”), but it is also the hardest to prove.

The microbiome change could be a *consequence* — the disease (or its treatment) alters the microbiome, rather than the other way around. A patient with Crohn’s disease who is experiencing severe intestinal inflammation will have a disrupted microbiome, but the disruption may simply reflect the inflamed, hostile environment in their gut rather than any causal role.

Or the microbiome change could be a *bystander* — both the disease and the microbiome shift are driven by a shared underlying factor (like diet, medication, or lifestyle), and they have no direct relationship with each other at all. This is the stork scenario, and it is far more common than most people realise.

The gold standard for establishing causality — the randomised controlled trial — is difficult to apply in microbiome science, because you cannot randomly assign people to have different microbiomes. Gnotobiotic mouse experiments (section 4.7) provide a partial answer, but with the caveats we have already discussed. Mendelian randomisation, a statistical technique that uses genetic variants as natural “experiments,” is beginning to be applied to microbiome data, but the approach is in its early stages.

Throughout this book, we will try to be honest about which of these categories each finding falls into. When the evidence supports a causal role for the microbiome, we will say so. When the evidence shows only an association, we will say that too — because knowing the difference between “these two things occur together” and “this thing causes that thing” is not a pedantic detail. It is the difference between understanding and illusion.

Compositional data and statistical traps. Microbiome data is *compositional* — it tells you the relative abundance of each organism in a sample, not the absolute abundance. If Species A increases from 10 per cent to 20 per cent of the total community, this could mean that Species A genuinely expanded, or it could mean that Species B (which was 50 per cent of the community) collapsed, causing everything else to increase as a proportion of the total without any actual change in abundance. Standard statistical methods designed for absolute counts can produce misleading results when applied to compositional data, and specialised statistical tools are needed.

Publication bias and hype. Microbiome science, like any rapidly growing field, is susceptible to publication bias — the tendency for positive, novel, or dramatic results to be published while negative or null results remain in file drawers. The Walter et al. critique of gnotobiotic studies, noted above, found that 95 per cent of published HMA mouse experiments reported positive results — a rate that is almost certainly inflated. Similarly, the media tends to amplify the most sensational microbiome findings (“Your gut bacteria control your mood!”) while ignoring the caveats, the effect sizes, and the distance between a mouse experiment and a proven human therapy.

This does not mean that the findings described in this book are unreliable. But it does mean that the reader should pay attention to the evidence grading that accompanies each claim. A finding that has been replicated in multiple large human cohorts is on different evidential footing from a finding that rests on a single mouse study. We will flag this distinction consistently throughout the book.

4.12 What the Tools Have Shown Us — and What Remains to Be Seen

The transformation in our ability to study microbial communities over the past two decades has been as dramatic as any technological revolution in the history of biology. In the space of twenty years, we have gone from a world in which most human-associated microorganisms were unknown and unknowable, to one in which we can identify every species in a sample, read every gene in their collective genomes, measure the metabolites they produce, track their activity in real time, and begin to understand — however imperfectly — how they interact with each other and with us.

But it is worth ending this chapter with a note of humility. For all the power of these tools, our understanding of the microbiome remains far more descriptive than mechanistic. We can describe what is there. We can observe what changes. We are much less able to explain *why* — to provide the kind of detailed, mechanistic understanding of cause and effect that we have for, say, the cardiovascular system or the cell cycle. The microbiome is a system of thousands of interacting species, producing thousands of metabolites, in constant dialogue with a host immune system of extraordinary complexity. Understanding this system at a mechanistic level is, arguably, one of the hardest problems in biology.

The tools described in this chapter are necessary but not sufficient. They have lifted the veil — shown us what was invisible — but the interpretation of what we see is still in its early stages. As we turn, in the chapters that follow, to the specific microbial communities of the body and the ways they communicate with us, we will repeatedly encounter this tension between the richness of the data and the limits of our understanding. The toolkit is powerful. The questions it has raised may take a generation to answer.

In Part II, we begin the guided tour — starting with the body's largest and most complex microbial habitat: the gut.

Chapter 4 — Key References

- [REF:staley1985] Staley, J.T. & Konopka, A. (1985). Measurement of in situ activities of non-photosynthetic microorganisms in aquatic and terrestrial habitats. *Annual Review of Microbiology*, 39, 321–346. DOI: 10.1146/annurev.mi.39.100185.001541
- [REF:woese1977] Woese, C.R. & Fox, G.E. (1977). Phylogenetic structure of the prokaryotic domain: the primary kingdoms. *Proceedings of the National Academy of Sciences*, 74(11), 5088–5090. DOI: 10.1073/pnas.74.11.5088
- [REF:pace1997] Pace, N.R. (1997). A molecular view of microbial diversity and the biosphere. *Science*, 276(5313), 734–740. DOI: 10.1126/science.276.5313.734
- [REF:hmp2012] The Human Microbiome Project Consortium. (2012). Structure, function and diversity of the healthy human microbiome. *Nature*, 486(7402), 207–214. DOI: 10.1038/nature11234
- [REF:qin2010] Qin, J. et al. (2010). A human gut microbial gene catalogue established by metagenomic sequencing. *Nature*, 464(7285), 59–65. DOI: 10.1038/nature08821
- [REF:wensel2022] Wensel, C.R., Pluznick, J.L., Salzberg, S.L. & Sears, C.L. (2022). Next-generation sequencing: insights to advance clinical investigations of the microbiome. *Journal of Clinical Investigation*, 132(7), e154944. DOI: 10.1172/JCI154944

- [REF:knight2018] Knight, R. et al. (2018). Best practices for analysing microbiomes. *Nature Reviews Microbiology*, 16(7), 410–422. DOI: 10.1038/s41579-018-0029-9
- [REF:salter2014] Salter, S.J. et al. (2014). Reagent and laboratory contamination can critically impact sequence-based microbiome analyses. *BMC Biology*, 12, 87. DOI: 10.1186/s12915-014-0087-z
- [REF:lagier2018] Lagier, J.-C. et al. (2018). Culturing the human microbiota and culturomics. *Nature Reviews Microbiology*, 16, 540–550. DOI: 10.1038/s41579-018-0041-0
- [REF:jones2015] Jones, M.B. et al. (2015). Library preparation methodology can influence genomic and functional predictions in human microbiome research. *Proceedings of the National Academy of Sciences*, 112(45), 14024–14029. DOI: 10.1073/pnas.1519288112
- [REF:walter2020] Walter, J., Armet, A.M., Finlay, B.B. & Shanahan, F. (2020). Establishing or exaggerating causality for the gut microbiome: lessons from human microbiota-associated rodents. *Cell*, 180(2), 221–232. DOI: 10.1016/j.cell.2019.12.025
- [REF:ley2006] Ley, R.E., Turnbaugh, P.J., Klein, S. & Gordon, J.I. (2006). Microbial ecology: human gut microbes associated with obesity. *Nature*, 444(7122), 1022–1023. DOI: 10.1038/4441022a
- [REF:dandreano2021] D’Andreano, S. et al. (2021). Rapid real-time identification of fungi up to the species level with long amplicon nanopore sequencing. *Biology Methods and Protocols*, 6(1), bpaa026. DOI: 10.1093/biomethods/bpaa026
- [REF:mitsuhashi2017] Mitsuhashi, S. et al. (2017). A portable system for rapid bacterial composition analysis using a nanopore-based sequencer and laptop computer. *Scientific Reports*, 7, 5657. DOI: 10.1038/s41598-017-05772-5
- [REF:wang2016] Wang, J. & Jia, H. (2016). Metagenome-wide association studies: fine-mining the microbiome. *Nature Reviews Microbiology*, 14(8), 508–522. DOI: 10.1038/nr-micro.2016.83

Additional references to be verified and expanded during manuscript preparation.

Chapter 5: The Gut — Our Largest Microbial Organ

In the first four chapters, we established a new way of thinking about the human body — not as a solitary organism but as an ecosystem, shaped by microbial partnership across evolutionary time, seeded at birth, and studied through increasingly powerful molecular tools. Now it is time to visit the ecosystem itself.

We begin where most of the microbes are: the gut.

If the human body is a continent, the gastrointestinal tract is its most densely populated territory. From the acid bath of the stomach to the dark, oxygen-starved reaches of the colon, the alimentary canal presents a series of radically different environments — each with its own climate, its own chemistry, and its own characteristic community of microbial inhabitants. More bacterial cells reside in your large intestine than in any other habitat on or inside your body. More microbial genes are expressed there than in all other body sites combined. And more of the dialogue between you and your microbiome — the chemical conversation that shapes your immunity, your metabolism, your mood, and your risk of disease — takes place across the gut lining than anywhere else.

This chapter is a guided tour of that territory. We will work our way down the digestive tract, habitat by habitat, and then examine the structures and processes that make the gut microbiome so consequential for human health: the mucus layer that keeps microbes in their place, the short-chain fatty acids that function as a currency of cooperation between bacteria and host, the enteric nervous system that wires the gut into the brain, the contested concept of enterotypes, and the ecological principle of colonisation resistance — the idea that a diverse, established community is its own best defence against invaders.

5.1 Anatomy of a Habitat: Stomach, Small Intestine, Colon — Different Worlds

The gastrointestinal tract is a single continuous tube, roughly nine metres long from mouth to anus. But calling it “a tube” is about as informative as calling Australia “an island.” The environments along its length differ as dramatically as desert differs from rainforest, and the microbial communities that inhabit each region reflect those differences with precision.

Think of it as a river system. Near the source, the water moves fast, conditions are harsh, and only specially adapted organisms survive. Further downstream, the current slows, nutrients accumulate, and life becomes abundant and diverse. The gut works the same way — except that the “harsh source” is the stomach, the “fast-flowing middle reaches” are the small intestine, and the “slow, nutrient-rich estuary” is the colon [REF:tropini2017].

The stomach is the least hospitable microbial habitat in the body. Gastric acid maintains a pH between 1.5 and 3.5 — acidic enough to dissolve metal. This is not an accident. The stomach’s acidity is, among other things, a defence: a chemical barrier that kills the vast majority of microorganisms swallowed with food and saliva. Estimates suggest that you swallow roughly a trillion bacteria per day from your oral cavity alone, and the stomach destroys most of them before they reach the intestine [REF:sender2016].

But “most” is not “all.” A small number of species have evolved to tolerate or even exploit this extreme environment. The most famous is *Helicobacter pylori*, a spiral-shaped bacterium that burrows into the gastric mucus layer and neutralises the acid in its immediate vicinity by producing the enzyme urease, which converts urea to ammonia — a base. *H. pylori* colonises roughly half the world’s population and, in most people, causes no symptoms at all. In a minority, it drives chronic gastritis, peptic ulcers, and gastric cancer — a story we will return to in Chapter 28. Beyond *H. pylori*, the healthy stomach harbours a sparse community dominated by *Streptococcus*, *Prevotella*, and *Veillonella* species, at densities of perhaps 10^1 to 10^3 bacteria per gram of content — a thousand to a million times less than the colon [REF:tropini2017].

The small intestine — duodenum, jejunum, and ileum — is where digestion and nutrient absorption happen. It is a challenging habitat for microbes, though for different reasons than the stomach. Transit time is relatively rapid: food residue moves through the six-metre length of the small intestine in three to five hours, which gives bacteria limited time to establish themselves. Bile salts, secreted into the duodenum from the liver and gallbladder, act as potent antimicrobials — they are detergent-like molecules that disrupt bacterial cell membranes. And the small intestine is bathed in digestive enzymes (proteases, lipases, amylases) that can damage microbial cells as readily as they digest food.

The result is a gradient. The duodenum and jejunum are relatively sparse — perhaps 10^3 to 10^4 bacteria per millilitre of content, dominated by fast-growing, oxygen-tolerant species like *Streptococcus* and *Lactobacillus*. By the time you reach the ileum — the final section of the small intestine, just upstream of the valve that separates it from the colon — microbial density has increased to 10^7 to 10^8 per millilitre, and the community begins to resemble a transitional zone between the sparse, oxygen-tolerant upstream communities and the dense, anaerobic downstream ones.

The ileum is also where you find one of the most immunologically important structures in the gut: the Peyer’s patches — organised clusters of immune tissue embedded in the intestinal wall that sample the microbial contents flowing past. We will explore their role in immune education in Chapter 12. For now, the key point is that the small intestine is not merely a transit zone. It is an active interface between microbes and the immune system, despite its relatively low microbial density.

The colon is where everything changes.

Past the ileocaecal valve — the one-way gate between the small and large intestine — conditions shift dramatically. Transit slows to twelve to thirty-six hours. Bile salts have been largely reabsorbed. Oxygen levels plummet. And into this slow, dark, anaerobic environment, microbes pour.

The colon is, by almost any measure, the most densely populated microbial habitat on Earth. Bacterial density reaches 10^{11} to 10^{12} cells per gram of content — roughly one hundred billion to one trillion bacteria in every gram of colonic material [REF:sender2016]. To put this in perspective: a single gram of your colon contents contains more bacterial cells than there are stars in the Milky Way.

This extraordinary density is not an accident of anatomy. It is the product of a carefully maintained environment. The colon provides what microbes need: warmth (a steady 37°C), a constant supply of substrates (the dietary fibre, mucus glycoproteins, and other complex carbohydrates that survived digestion in the small intestine), and critically, the near-absence of oxygen. Most of the dominant colonic bacteria are *obligate anaerobes* — organisms that not only do not need oxygen but are poisoned by it. The colon’s low oxygen tension is actively maintained by the metabolic activity of

the epithelial cells lining the gut wall, which consume oxygen so avidly that they create an almost anoxic environment in the lumen [REF:tropini2017].

There is also a gradient within the colon itself. The proximal colon (the right side, beginning with the caecum) receives the most substrate — the freshly arriving fibre and residues from the small intestine — and here, fermentation is most active, pH is lower (around 5.5 to 6.0), and microbial metabolism is at its most energetic. The distal colon (the left side, ending at the rectum) receives what is left after the proximal community has had first access. Substrate availability declines, pH rises toward neutral, and the microbial community shifts accordingly. Some species that thrive in the substrate-rich proximal colon are less abundant distally, and vice versa.

This anatomy matters because it means that the gut is not one habitat but many. The microbial community of the stomach is as different from that of the colon as a desert community is from a tropical forest. And even within the colon, there are spatial distinctions — between the lumen (the central space) and the mucosa (the tissue lining the wall), between the proximal and distal segments, and between the crypts (the microscopic indentations in the mucosal surface, which harbour specialised communities) and the exposed epithelial surface. A stool sample — the most commonly collected specimen in microbiome research — captures the luminal community of the distal colon reasonably well, but it tells you relatively little about what is happening at the mucosal surface, in the small intestine, or in the proximal colon. This is one of the important methodological caveats that readers should keep in mind throughout the chapters that follow.

5.2 The Dominant Phyla and What Drives Community Structure

If you could take a census of every bacterium in a healthy adult colon, the results would be remarkably consistent across individuals, at least at the broadest level of classification.

Two bacterial phyla dominate. *Bacillota* (formerly known as Firmicutes) and *Bacteroidota* (formerly Bacteroidetes) together typically account for more than 90 per cent of all bacteria in the gut. This has been confirmed in every large-scale survey since the Human Microbiome Project, and it holds true across geography, ethnicity, and diet — though the proportions vary [REF:hmp2012].

A word on the names. The phylum-level taxonomy of bacteria was revised in 2021, and many of the familiar names that have appeared in popular science writing for two decades were replaced. Firmicutes became *Bacillota*. Bacteroidetes became *Bacteroidota*. Actinobacteria became *Actinomycetota*. Proteobacteria became *Pseudomonadota*. The old names remain so deeply embedded in the literature — and in the reader's Google search results — that we will use both: the current name first, followed by the former name in parentheses on first mention. After that, we use the current names.

Bacillota is an enormous phylum containing hundreds of genera. In the gut, the key players include *Faecalibacterium* (especially *F. prausnitzii*, one of the most abundant single species in the healthy human colon and a major producer of the short-chain fatty acid butyrate), *Roseburia*, *Eubacterium*, *Ruminococcus*, *Clostridium* (a large and diverse genus that includes both beneficial fibre fermenters and dangerous pathogens like *C. difficile*), and *Lactobacillus*. Many *Bacillota* are specialist fibre degraders — they possess the enzymatic machinery to break down complex plant polysaccharides that human enzymes cannot touch.

Bacteroidota is dominated in the gut by the genus *Bacteroides* — versatile metabolic generalists

that can switch between different carbohydrate substrates depending on what is available. *Bacteroides thetaiotaomicron*, one of the most studied species in microbiome science, possesses over 260 genes dedicated to breaking down complex carbohydrates — far more than the human genome encodes. The *Prevotella* genus, also within Bacteroidota, tends to be more abundant in individuals consuming high-fibre, plant-rich diets, while *Bacteroides* is more associated with Western diets rich in protein and animal fat [REF:wu2011].

Beyond the two dominant phyla, several others make important contributions. *Actinomycetota* (formerly Actinobacteria) includes the *Bifidobacterium* genus — major players in the infant gut and prominent ingredients in commercial probiotics, though their relative abundance declines in adulthood. *Pseudomonadota* (formerly Proteobacteria) is a phylum that includes *Escherichia coli* and other Gram-negative species; in a healthy gut, Pseudomonadota represent a small minority (typically less than 5 per cent), but their expansion — a phenomenon sometimes called a “bloom” — is often associated with inflammation, antibiotic use, or disease, and has been proposed as a diagnostic signature of microbial disruption, or *dysbiosis* [REF:rinninella2019].

Verrucomicrobiota is represented in the gut almost entirely by a single species: *Akkermansia muciniphila*, a mucus-degrading specialist that lives in the mucus layer lining the colon and has attracted intense interest for its associations with metabolic health, lean body mass, and improved responses to cancer immunotherapy. We will encounter *Akkermansia* again in several later chapters.

What drives the relative proportions of these groups? The short answer is: diet, more than almost anything else. The landmark study by Gary Wu and colleagues at the University of Pennsylvania, published in *Science* in 2011, demonstrated that long-term dietary patterns were strongly correlated with gut community composition. Individuals whose habitual diets were high in protein and animal fat had gut communities enriched in *Bacteroides*, while those eating high-fibre, plant-based diets were enriched in *Prevotella* [REF:wu2011]. This dietary signal is so strong that it can override geography, genetics, and ethnicity — populations in rural Africa eating traditional high-fibre diets have gut communities that look more like those of rural South Americans eating similar diets than they do like those of urban Africans eating Westernised diets.

But diet is not the only driver. Host genetics plays a modest role — twin studies have shown that certain bacterial taxa are more heritable than others, with *Christensenellaceae* being the most heritable family identified so far. Medications, particularly antibiotics and proton pump inhibitors, can profoundly reshape the community (topics we will explore in Chapters 19 and 22). Age matters: the infant gut is dominated by *Bifidobacterium*, the adult gut by Bacillota and Bacteroidota, and the elderly gut often shows declining diversity and rising Pseudomonadota. Geography, birth mode, breastfeeding history, household contacts, pet ownership — all leave detectable signatures in the gut microbiome. The community you carry is the product of everything that has happened to you, filtered through the selective environment of your own particular gut.

5.3 The Mucus Layer: A Negotiated Border

The gut faces a paradox. It must absorb nutrients from its contents — which requires an enormous surface area of thin, permeable tissue — while simultaneously keeping trillions of bacteria from crossing that tissue and entering the bloodstream. If the barrier fails, the result is not a minor inconvenience. Bacteria or their products reaching the blood can trigger systemic inflammation, sepsis, and death. The gut, in other words, must be simultaneously open and closed: open to food,

closed to microbes.

The primary structure that solves this paradox is not the epithelial cells themselves, but the mucus layer that coats them.

Intestinal mucus is a hydrated gel composed primarily of a single, enormous glycoprotein called MUC2 — one of the largest molecules the human body produces, with a molecular weight of over five million daltons. MUC2 is secreted by goblet cells, specialised epithelial cells scattered among the absorptive cells that line the intestinal surface. In the colon, the mucus that these cells produce is organised into two distinct layers — an arrangement first definitively demonstrated by Gunnar Hansson's group in Gothenburg, Sweden, in a landmark 2008 study [REF:johansson2008].

The **outer mucus layer** is loose, expanded, and teeming with bacteria. It is, in effect, a microbial habitat — a structured environment where bacteria find attachment sites, nutrients (the sugar chains of the mucin glycoproteins themselves are a food source), and protection from the flow of intestinal contents. Many of the bacteria that live in the outer mucus layer are different from those floating freely in the lumen. Mucus specialists like *Akkermansia muciniphila* thrive here, using enzymes to trim the sugar chains from mucin molecules and feeding on the released sugars.

The **inner mucus layer** is dense, firmly attached to the epithelial surface, and — critically — almost entirely devoid of bacteria. Johansson and colleagues showed this by using fluorescent staining to visualise both the mucus and the bacteria simultaneously: the inner layer appeared as a clear, bacteria-free zone, like a no-man's-land between the microbial masses in the outer layer and the delicate epithelial cells beneath [REF:johansson2008]. The inner layer is approximately 50 micrometres thick in mice (somewhat variable in humans) and is continuously secreted and renewed. Its impermeability to bacteria is not merely physical — it also contains antimicrobial peptides, secretory IgA antibodies, and other immune molecules that actively kill or neutralise any microbes that penetrate too deeply.

The analogy that works best here is a border zone. Imagine a country that trades heavily with its neighbour but maintains a strictly controlled border. The outer mucus layer is the bustling border town — full of activity, commerce, and foreign nationals going about their business. The inner mucus layer is the fortified border wall itself — patrolled, reinforced, and designed to prevent unauthorised crossings. The epithelial cells behind it are the homeland.

This two-layer system is not static. It requires constant maintenance. Goblet cells continuously secrete new mucus, which is added to the inner layer. The inner layer gradually converts into the outer layer as it expands and becomes colonised by bacteria. And the bacteria in the outer layer continuously degrade the mucus they inhabit — consuming the sugar chains as food, which thins the layer from the outside even as it is replenished from the inside. The result is a dynamic equilibrium: mucus is produced, colonised, degraded, and replaced in a continuous cycle.

What happens when this equilibrium breaks down? One striking answer came from a 2016 study by Mahesh Desai and colleagues, published in *Cell* [REF:desai2016]. Using gnotobiotic mice colonised with a defined community of human gut bacteria, the team showed that when dietary fibre was removed from the animals' food, the bacteria that normally fermented fibre switched to the only other available complex carbohydrate: the mucus layer itself. Deprived of their preferred substrate, the microbes ate the border wall. The mucus layer thinned dramatically, and when the researchers then challenged the mice with *Citrobacter rodentium* (a pathogen used as a model for human enteric infections), the fibre-deprived mice developed far more severe infections than controls fed a fibre-rich diet.

The implication is both elegant and unsettling. Dietary fibre is not merely “good for digestion” in some vague, cereal-box sense. It is the substrate that keeps the mucus layer intact. When fibre runs out, the bacteria eat the barrier — and the barrier’s degradation makes the host more vulnerable to infection. The mucus layer is not just a passive shield. It is an actively maintained structure whose integrity depends on a continuous negotiation between the host (which produces it) and the microbiota (which both inhabit and consume it). Both parties must hold up their end of the deal.

The 2016 Birchenough study from Hansson’s group added another layer to this picture: the discovery of “sentinel goblet cells” positioned at the entrances to colonic crypts. These specialised cells detect bacterial contact through the NLRP6 inflammasome and respond by triggering a burst of mucus secretion — effectively sealing the crypt entrance when bacteria approach too closely [REF:birchenough2016]. It is an active defence, triggered by microbial incursion, that protects the stem cells deep in the crypts from bacterial contact.

The mucus layer, then, is far more than a simple physical barrier. It is a dynamic, multi-layered structure that serves simultaneously as habitat, food source, immune interface, and border wall. Its maintenance is one of the most important — and most underappreciated — functions of the gut ecosystem.

5.4 Short-Chain Fatty Acids and Microbial Metabolism — The Currency of Co-operation

If the mucus layer is the border between host and microbe, short-chain fatty acids are the currency that flows across it. They are, arguably, the single most important class of molecules produced by the gut microbiome — the primary mechanism through which dietary fibre is converted into benefits for the host.

The chemistry is straightforward. Humans lack the enzymes to break down most complex plant polysaccharides — the cellulose, hemicellulose, resistant starch, inulin, pectins, and beta-glucans collectively known as *dietary fibre* or, in microbiome research, *microbiota-accessible carbohydrates* (MACs). These molecules pass through the stomach and small intestine essentially unchanged. When they reach the colon, they are fermented by anaerobic bacteria — broken down through a series of metabolic steps into smaller molecules. The principal end products of this fermentation are three short-chain fatty acids: **acetate** (two carbon atoms), **propionate** (three carbons), and **butyrate** (four carbons). Together, these three molecules are produced in enormous quantities — roughly 500 to 600 millimoles per day in a person consuming a fibre-rich diet — and they are absorbed rapidly across the colonic epithelium into the host [REF:cummings1987].

The classic measurements of SCFA concentrations in the human gut were published by John Cummings and colleagues in 1987, based on direct sampling of portal and peripheral blood as well as colonic contents — an invasive study that has never been repeated at comparable scale. They established that the proximal colon is the primary site of SCFA production, that acetate is the most abundant (comprising roughly 60 per cent of total SCFAs), followed by propionate (~25 per cent) and butyrate (~15 per cent), and that the majority of SCFAs produced are absorbed before reaching the distal colon [REF:cummings1987].

What makes SCFAs so important is not merely that they are produced in large quantities, but that they affect the host at almost every level — local, systemic, and even neurological [REF:koh2016].

Locally, butyrate is the primary fuel of the colonic epithelium. The cells that line the colon — the colonocytes — derive roughly 60 to 70 per cent of their energy from butyrate oxidation. This is a remarkable metabolic arrangement: the host cells that are in closest contact with bacteria are powered by a molecule that only bacteria can produce. If butyrate production falls — because of antibiotic use, a low-fibre diet, or a disrupted microbiome — colonocytes are energy-starved, the epithelial barrier weakens, and the mucus layer thins. The dependency is real and consequential.

Butyrate also has a direct role in maintaining the low-oxygen environment that the anaerobic bacteria need to survive. When colonocytes oxidise butyrate, they consume oxygen — and this oxygen consumption helps maintain the steep oxygen gradient between the oxygenated tissue and the nearly anoxic lumen. In a sense, the bacteria produce the fuel that keeps the host cells consuming the oxygen that keeps the environment suitable for the bacteria. It is a feedback loop of mutual dependency.

Systemically, SCFAs enter the bloodstream and influence distant organs. Propionate is largely taken up by the liver, where it feeds into gluconeogenesis (the production of glucose) and helps regulate cholesterol synthesis. Acetate reaches the systemic circulation in significant quantities and has been shown to influence appetite regulation, fat storage, and immune function. Both propionate and butyrate activate a family of cell-surface receptors called free fatty acid receptors (FFAR2 and FFAR3, also known as GPR43 and GPR41), which are expressed on immune cells, adipocytes, enteroendocrine cells, and neurons — providing a direct molecular link between microbial metabolism and host physiology [REF:koh2016].

Immunologically, butyrate and propionate are potent regulators of inflammation. Both molecules inhibit histone deacetylases (HDACs) — enzymes that control gene expression by modifying the proteins around which DNA is wound. By inhibiting HDACs, butyrate promotes the differentiation of regulatory T cells (Tregs) — the immune cells responsible for suppressing excessive inflammation and maintaining tolerance. This is one of the key molecular mechanisms linking a fibre-rich diet to reduced inflammation and a lower risk of autoimmune disease: dietary fibre feeds butyrate-producing bacteria; butyrate promotes Treg differentiation; Tregs keep inflammation in check. Remove any link in this chain and the system falters.

The analogy that captures the SCFA relationship best is that of a local economy. The bacteria are manufacturers, producing SCFAs from raw materials (fibre) that the host provides through diet. The host is both the supplier of raw materials and the consumer of the finished product. The economy works when both sides participate: the host eats fibre, the bacteria produce SCFAs, the host's cells use the SCFAs for energy and immune regulation, and in return the host maintains the warm, anaerobic, substrate-rich environment the bacteria need. When the supply of raw materials dries up — when the host stops eating fibre — the economy collapses. The manufacturers switch to degrading the infrastructure itself (the mucus layer), the energy supply to colonocytes drops, the barrier weakens, and inflammation rises.

This is why dietary fibre is not merely a “health food” but the foundational input of the gut ecosystem. We will return to this theme at length in Chapter 20.

5.5 The Enteric Nervous System: A “Second Brain” Wired Into the Microbiome

Embedded in the walls of the gastrointestinal tract is a network of roughly 500 million neurons — more than in the spinal cord, and more than in any other organ outside the brain. This network, the *enteric nervous system* (ENS), is sometimes called the “second brain,” and while the label is imprecise (the ENS does not think, plan, or compose poetry), it captures something real: the gut has its own autonomous nervous system, capable of coordinating digestion, motility, blood flow, and immune responses without any input from the brain [REF:furness2012].

The ENS is organised into two main networks, or *plexuses*. The *myenteric plexus* (Auerbach’s plexus) lies between the two muscle layers of the gut wall and controls peristalsis — the rhythmic contractions that move food along the tract. The *submucosal plexus* (Meissner’s plexus) lies closer to the luminal surface and regulates secretion, absorption, and local blood flow. Together, these plexuses form an integrated circuit that can operate the entire digestive process independently — which is why a segment of gut removed from the body and placed in a nutrient bath will continue to contract rhythmically on its own.

What makes the ENS relevant to this chapter is its proximity to, and interaction with, the gut microbiome. The neurons of the submucosal plexus lie just micrometres from the epithelial surface — and therefore just micrometres from the microbial community on the other side of that surface. The microbiome influences the ENS through multiple channels.

The most striking example involves serotonin. Roughly 90 per cent of the body’s total serotonin — a neurotransmitter best known for its role in mood regulation — is produced not in the brain but in the gut, by specialised epithelial cells called enterochromaffin cells. In 2015, Elaine Hsiao’s group at UCLA published a landmark study in *Cell* demonstrating that gut bacteria — specifically spore-forming bacteria, primarily Clostridia — regulate serotonin production by these cells [REF:yano2015]. Germ-free mice, which lack a gut microbiome, produce roughly 60 per cent less serotonin in their colon than conventionally colonised mice. When the researchers colonised germ-free mice with spore-forming bacteria, serotonin production was restored. The mechanism involves bacterial metabolites — including SCFAs and secondary bile acids — that signal to enterochromaffin cells and upregulate the expression of tryptophan hydroxylase 1, the rate-limiting enzyme in serotonin synthesis.

Gut-derived serotonin does not cross the blood-brain barrier, so it does not directly alter mood in the way that brain serotonin does. But it plays critical roles in regulating gut motility (serotonin is a major driver of peristalsis), secretion, visceral pain perception, and inflammation. It also enters the bloodstream and influences platelet function, bone metabolism, and liver regeneration. The gut microbiome, through its influence on serotonin production, is quietly modulating processes throughout the body.

The ENS is also connected to the brain via the *vagus nerve* — the longest cranial nerve in the body, running from the brainstem to the abdomen and carrying both sensory (afferent) and motor (efferent) signals in both directions. About 80 per cent of the vagus nerve’s fibres are afferent — they carry information *from* the gut *to* the brain, not the other way around. This means that the gut is, in a quantitative sense, primarily an information sender rather than an information receiver. Microbial metabolites, immune signals, and ENS activity are continuously transmitted to the brain via the vagus nerve, where they influence mood, appetite, stress responses, and even memory.

We will explore the gut-brain axis in much greater detail in Chapter 13. For now, the essential point

is that the gut is not merely a digestive organ with a microbial side-show. It is a sensory organ — a vast neural network immersed in a microbial ecosystem, continuously sampling the chemical environment and relaying that information to the central nervous system. The microbes in your gut are not isolated from your nervous system. They are wired into it.

5.6 Enterotypes — Do They Exist, and Do They Matter?

In 2011, a large European consortium led by Peer Bork published a paper in *Nature* that proposed something provocative: that the human gut microbiome could be classified into a small number of distinct community types — which they called *enterotypes* — analogous to blood types [REF:arumugam2011].

Analysing metagenomic data from 39 individuals from three continents, the team identified three clusters, each characterised by the dominance of a different genus: *Bacteroides* (Enterotype 1), *Prevotella* (Enterotype 2), and *Ruminococcus* (Enterotype 3). The idea was immediately attractive. If people's gut communities fell into a few discrete categories, it would simplify both research and clinical practice enormously. You could stratify patients by enterotype. You could predict responses to diet or drugs. You could design targeted interventions.

The paper was widely cited and widely discussed. It was also widely challenged.

The first problem was statistical. Subsequent analyses with larger datasets had difficulty reproducing the three clean clusters. The *Ruminococcus* enterotype proved especially unstable — in many studies, it merged with the *Bacteroides* cluster or disappeared entirely. By 2014, Dan Knights and colleagues argued in *Cell Host & Microbe* that the data were better described as a continuous gradient rather than discrete types — that people's gut communities varied along smooth dimensions rather than falling into bins [REF:knights2014]. The blood type analogy, they suggested, was misleading. A better analogy might be height: people vary continuously, and while you can draw arbitrary lines to create categories (short, medium, tall), the categories do not reflect a biological discontinuity.

The counterargument came in 2018, when Costea and colleagues (including Bork's group) published a large-scale reanalysis in *Nature Microbiology*, incorporating data from nearly 4,000 samples [REF:costea2018]. They conceded that the original three-enterotype model was an oversimplification, but argued that at least two robust community types — one dominated by *Bacteroides*, the other by *Prevotella* — were consistently identifiable across datasets. These were not discrete bins with sharp boundaries, but they were genuine density peaks in the landscape of community composition — hills on a map, even if the terrain between them was continuous.

What drives the *Bacteroides/Prevotella* split? Diet appears to be the primary factor. The Wu et al. *Science* study, also from 2011, showed that long-term dietary habits were strongly associated with enterotype: habitual high-protein, high-fat diets correlated with *Bacteroides* dominance, while high-fibre, plant-rich diets correlated with *Prevotella* [REF:wu2011]. This fits with what we know about the metabolic capabilities of these genera: *Bacteroides* species are versatile metabolisers of proteins and animal-derived substrates, while *Prevotella* species specialise in degrading plant polysaccharides.

Does the distinction matter clinically? The evidence is suggestive but not yet definitive. Some studies have found that individuals with *Prevotella*-dominant communities respond differently to

dietary interventions than those with *Bacteroides*-dominant communities — losing more weight on fibre-rich diets, for example. Others have linked *Prevotella* dominance to different patterns of immune activation. But the field has not yet reached the point where enterotyping (or whatever the continuous analogue might be called) has practical clinical utility.

The current consensus might be summarised as follows: the gut microbiome is not as neatly classifiable as blood types, but neither is it a featureless continuum. There are real compositional patterns, driven primarily by diet, that cluster people's microbiomes into recognisable — if fuzzy — zones. Whether these zones will eventually prove clinically useful remains to be seen. For the moment, the enterotype concept is best understood as an imperfect but informative way of describing the landscape of human gut community variation — a topographic map that shows the hills, even if the borders between them are drawn in pencil rather than ink.

5.7 Colonisation Resistance — Why a Full House Is Hard to Invade

We end this chapter with a concept that will recur throughout the rest of the book, one that is essential to understanding both how the gut microbiome protects us and why well-intentioned interventions sometimes backfire.

In 1954, Marjorie Bohnhoff and colleagues at the University of Chicago performed an experiment that would prove prescient. They treated mice with the antibiotic streptomycin and then exposed them to *Salmonella* bacteria. The antibiotic-treated mice developed severe infections at doses that barely affected untreated animals. Something about the normal intestinal flora, Bohnhoff concluded, was protecting the host from pathogen colonisation [REF:bohnhoff1954].

The concept was given a name in 1971 by Dirk van der Waaij and colleagues, who called it *colonisation resistance*: the ability of an established microbial community to prevent the establishment of new organisms, including pathogens [REF:vanderwaaij1971]. It is one of the most important ecological services the gut microbiome provides — and one of the least appreciated by the general public.

How does colonisation resistance work? The mechanisms are multiple and layered, operating through at least four distinct pathways [REF:buffie2013].

First, nutrient competition. An established gut community has already occupied the available metabolic niches — the bacteria present are consuming the available substrates, and there is little left over for a newcomer. This is the ecological principle of *competitive exclusion*: in a stable environment, two species competing for exactly the same resource cannot coexist indefinitely — one will outcompete the other. A diverse, established community occupies so many niches simultaneously that an incoming species, no matter how well-adapted, struggles to find an unoccupied niche and gain a foothold. The pantry is not merely stocked — it is being raided by dozens of specialists who know exactly where everything is.

Second, direct antimicrobial production. Many commensal bacteria produce *bacteriocins* — small antimicrobial peptides that kill or inhibit closely related species — as well as organic acids and other toxic metabolites. These are not accidental by-products. They are weapons, deployed in the constant low-level warfare between bacterial species competing for territory and resources. The resident community, collectively, maintains a chemical environment that is hostile to newcomers.

Third, bile acid modification. Primary bile acids, synthesised by the liver and secreted into the

duodenum, are modified by gut bacteria into secondary bile acids through a process called deconjugation and 7 α -dehydroxylation. The secondary bile acids produced — particularly deoxycholic acid and lithocholic acid — are potent inhibitors of many pathogens, including *Clostridioides difficile*. When antibiotics eliminate the bile-acid-modifying bacteria, secondary bile acid concentrations plummet, and *C. difficile* — which can germinate its spores in the presence of primary bile acids but is inhibited by secondary bile acids — finds a suddenly permissive environment. This mechanism is one of the primary reasons that antibiotic use is the dominant risk factor for *C. difficile* infection [REF:buffie2013].

Fourth, immune priming. The resident microbiome continuously stimulates the gut immune system — promoting the secretion of IgA antibodies into the mucus layer, maintaining the production of antimicrobial peptides by Paneth cells (specialised epithelial cells in the small intestinal crypts), and keeping the immune system in a state of calibrated readiness. This is not a response to a specific threat. It is a tonic signal — a baseline level of immune activation maintained by the presence of commensal bacteria, which incidentally provides a first line of defence against any pathogen that attempts to breach the barrier [REF:ducarmon2019].

The analogy here is ecological, and it is worth stating plainly: the gut microbiome protects the host from infection in the same way that a healthy forest resists invasion by weeds. A mature forest — with its canopy blocking sunlight, its root systems occupying the soil, its established network of mycorrhizal fungi and nutrient cycles — is almost impossible for an invasive plant species to colonise. But clear-cut the forest (antibiotics), and suddenly the ground is bare, the soil is exposed, and any windblown seed has a chance. The first species to arrive after the disruption are not necessarily the “best” species — they are the fastest, the most opportunistic, the ones whose spores were already present and waiting. This is why *C. difficile* and *Candida* species bloom after antibiotic use: not because antibiotics create these organisms, but because antibiotics remove the community that was keeping them in check.

This ecological principle — that a diverse, established community resists invasion by newcomers — has a corollary that the reader should hold in mind as we move through the rest of this book. It applies not only to pathogens. It applies to *any* newcomer, including organisms that we might consider “beneficial.” An established gut community does not distinguish between a pathogenic *Salmonella* and a probiotic *Lactobacillus* arriving in a capsule. Both are outsiders. Both face the same competitive barriers. Both must find an unoccupied niche or displace a resident to establish themselves.

This is not an argument against probiotics — that discussion will come in Chapter 23, after we have laid considerably more groundwork. It is simply a statement of ecological principle: colonisation resistance is not selective. A healthy community resists all newcomers, not just the ones we wish it would resist. Understanding this principle is essential to understanding both why probiotics often fail to colonise the gut (they are outcompeted by the residents) and why the vulnerable moments — after antibiotics, in early infancy, during illness — are both the most dangerous and the most therapeutically interesting.

The gut is not one place. It is a series of habitats, grading from the acidic desolation of the stomach through the rapid-transit corridor of the small intestine to the dense, teeming, anaerobic metropolis of the colon. Its microbial community is shaped by diet, maintained by mucus, fuelled by fibre, wired into the nervous system, and protected by the very ecological principles that make it so resilient.

In the next chapter, we leave the interior and move to the surface. The skin — your body’s largest organ by area — is home to a microbial community every bit as complex, every bit as site-specific, and every bit as consequential as the one we have just explored. But the rules are different. On the skin, the challenges are not acid and bile but UV radiation, desiccation, and an immune system that treats the surface as a battleground. It is, quite literally, a different world.

Chapter 5 — Key References

- [REF:tropini2017] Tropini, C., Earle, K.A., Huang, K.C. & Sonnenburg, J.L. (2017). The Gut Microbiome: Connecting Spatial Organization to Function. *Cell Host & Microbe*, 21(4), 433–442. DOI: 10.1016/j.chom.2017.03.010
- [REF:sender2016] Sender, R., Fuchs, S. & Milo, R. (2016). Revised estimates of the number of human and bacteria cells in the body. *PLOS Biology*, 14(8), e1002533. DOI: 10.1371/journal.pbio.1002533
- [REF:hmp2012] The Human Microbiome Project Consortium. (2012). Structure, function and diversity of the healthy human microbiome. *Nature*, 486(7402), 207–214. DOI: 10.1038/nature11234
- [REF:johansson2008] Johansson, M.E.V., Phillipson, M., Petersson, J., Velcich, A., Holm, L. & Hansson, G.C. (2008). The inner of the two Muc2 mucin-dependent mucus layers in colon is devoid of bacteria. *Proceedings of the National Academy of Sciences*, 105(39), 15064–15069. DOI: 10.1073/pnas.0803124105
- [REF:desai2016] Desai, M.S., Seekatz, A.M., Koropatkin, N.M., Kamada, N., Hickey, C.A., Wolter, M., ... & Martens, E.C. (2016). A Dietary Fiber-Deprived Gut Microbiota Degrades the Colonic Mucus Barrier and Enhances Pathogen Susceptibility. *Cell*, 167(5), 1339–1353. DOI: 10.1016/j.cell.2016.10.043
- [REF:birchenough2016] Birchenough, G.M.H., Nyström, E.E.L., Johansson, M.E.V. & Hansson, G.C. (2016). A sentinel goblet cell guards the colonic crypt by triggering Nlrp6-dependent Muc2 secretion. *Science*, 352(6293), 1535–1542. DOI: 10.1126/science.aaf7419
- [REF:hansson2020] Hansson, G.C. (2020). Mucins and the Microbiome. *Annual Review of Biochemistry*, 89, 769–793. DOI: 10.1146/annurev-biochem-011520-105053
- [REF:cummings1987] Cummings, J.H., Pomare, E.W., Branch, W.J., Naylor, C.P. & Macfarlane, G.T. (1987). Short chain fatty acids in human large intestine, portal, hepatic and venous blood. *Gut*, 28(10), 1221–1227. DOI: 10.1136/gut.28.10.1221
- [REF:koh2016] Koh, A., De Vadder, F., Kovatcheva-Datchary, P. & Bäckhed, F. (2016). From Dietary Fiber to Host Physiology: Short-Chain Fatty Acids as Key Bacterial Metabolites. *Cell*, 165(6), 1332–1345. DOI: 10.1016/j.cell.2016.05.041
- [REF:morrison2016] Morrison, D.J. & Preston, T. (2016). Formation of short chain fatty acids by the gut microbiota and their impact on human metabolism. *Gut Microbes*, 7(3), 189–200. DOI: 10.1080/19490976.2015.1134082
- [REF:furness2012] Furness, J.B. (2012). The enteric nervous system and neurogastroenterology. *Nature Reviews Gastroenterology & Hepatology*, 9(5), 286–294. DOI: 10.1038/nrgastro.2012.32

- [REF:yano2015] Yano, J.M., Yu, K., Donaldson, G.P., Shastri, G.G., Ann, P., Ma, L., ... & Hsiao, E.Y. (2015). Indigenous bacteria from the gut microbiota regulate host serotonin biosynthesis. *Cell*, 161(2), 264–276. DOI: 10.1016/j.cell.2015.02.047
- [REF:arumugam2011] Arumugam, M., Raes, J., Pelletier, E., Le Paslier, D., Yamada, T., Mende, D.R., ... & Bork, P. (2011). Enterotypes of the human gut microbiome. *Nature*, 473(7346), 174–180. DOI: 10.1038/nature09944
- [REF:knights2014] Knights, D., Ward, T.L., McKinlay, C.E., Miller, H., Gonzalez, A., McDonald, D. & Knight, R. (2014). Rethinking “enterotypes.” *Cell Host & Microbe*, 16(4), 433–437. DOI: 10.1016/j.chom.2014.09.013
- [REF:costea2018] Costea, P.I., Hildebrand, F., Arumugam, M., Bäckhed, F., Blaser, M.J., Bushman, F.D., ... & Bork, P. (2018). Enterotypes in the landscape of gut microbial community composition. *Nature Microbiology*, 3(1), 8–16. DOI: 10.1038/s41564-017-0072-8
- [REF:wu2011] Wu, G.D., Chen, J., Hoffmann, C., Bittinger, K., Chen, Y.-Y., Keilbaugh, S.A., ... & Lewis, J.D. (2011). Linking long-term dietary patterns with gut microbial enterotypes. *Science*, 334(6052), 105–108. DOI: 10.1126/science.1208344
- [REF:rinninella2019] Rinninella, E., Raoul, P., Cintoni, M., Franceschi, F., Miggianno, G.A.D., Gasbarrini, A. & Mele, M.C. (2019). What is the Healthy Gut Microbiota Composition? A Changing Ecosystem across Age, Environment, Diet, and Diseases. *Microorganisms*, 7(1), 14. DOI: 10.3390/microorganisms7010014
- [REF:bohnhoff1954] Bohnhoff, M., Drake, B.L. & Miller, C.P. (1954). Effect of streptomycin on susceptibility of intestinal tract to experimental Salmonella infection. *Proceedings of the Society for Experimental Biology and Medicine*, 86(1), 132–137.
- [REF:vanderwaaij1971] van der Waaij, D., Berghuis-de Vries, J.M. & Lekkerkerk-van der Wees, J.E.C. (1971). Colonization resistance of the digestive tract in conventional and antibiotic-treated mice. *Journal of Hygiene*, 69(3), 405–411. DOI: 10.1017/S0022172400021653
- [REF:buffie2013] Buffie, C.G. & Pamer, E.G. (2013). Microbiota-mediated colonization resistance against intestinal pathogens. *Nature Reviews Immunology*, 13(11), 790–801. DOI: 10.1038/nri3535
- [REF:ducarmon2019] Ducarmon, Q.R., Zwitterink, R.D., Hornung, B.V.H., van Schaik, W., Young, V.B. & Kuijper, E.J. (2019). Gut Microbiota and Colonization Resistance against Bacterial Enteric Infection. *Microbiology and Molecular Biology Reviews*, 83(3), e00007-19. DOI: 10.1128/MMBR.00007-19

Additional references to be verified and expanded during manuscript preparation.

Chapter 6: The Skin — Life on the Surface

In the previous chapter, we explored the gut — the body’s most densely populated microbial habitat, a warm, dark, nutrient-rich interior world. Now we step outside.

The skin is the human body’s largest organ, stretching roughly 1.8 square metres in an average adult and weighing about four kilograms. It is also, by any reckoning, the most exposed. Every surface you touch, every handshake, every gust of wind deposits microorganisms on your skin or carries them away. If the gut is a tropical rainforest — lush, sheltered, teeming with life — then the skin is a continent of extremes: scorching deserts, humid valleys, oily swamps, and cool, dry plains, all existing simultaneously on a single human body.

And yet, despite this exposure, despite the constant bombardment from the outside world, the skin is not simply overrun. It is selectively colonised. The communities that live on your skin are not random assemblages of whatever happened to land there. They are structured, site-specific, and remarkably stable — shaped by the physical and chemical properties of each patch of skin, by the immune system patrolling beneath, and by the competitive interactions among the microbes themselves.

This chapter is a tour of that outer landscape. We will explore how different regions of the skin support radically different microbial communities, why some microbes are permanent residents while others are fleeting visitors, how the skin’s own defences work in concert with its microbial inhabitants, and what happens when this partnership breaks down — in conditions from acne to eczema to chronic wounds. We will also meet *Malassezia*, the overlooked fungal genus that dominates the skin mycobiome, and close with a question that connects the skin’s microbial ecology to one of the most debated ideas in modern immunology: the hygiene hypothesis.

6.1 Deserts, Rainforests, and Caves: Ecological Niches Across the Skin

If someone asked you to describe the environment of your skin, you might offer a single answer: warm, slightly acidic, mostly dry. But that answer would be about as accurate as describing Australia’s climate as “hot.” The skin is not one environment. It is dozens, and the microbial communities that inhabit them reflect this diversity with extraordinary precision.

The landmark study that revealed this was published in *Science* in 2009 by Elizabeth Grice, Julia Segre, and colleagues at the National Institutes of Health. They sampled 20 distinct skin sites on healthy volunteers and sequenced the bacterial DNA they found, using the 16S ribosomal RNA gene as a molecular barcode. What emerged was a map of microbial geography as detailed and varied as any atlas of terrestrial ecosystems [REF:grice2009].

The key insight was that skin sites cluster into three broad habitat types, each defined by a physical property of the skin itself: **sebaceous** (oily) sites, **moist** sites, and **dry** sites. The microbial communities at sites sharing the same habitat type resemble each other far more than they resemble the communities at nearby sites of a different type. In other words, the microbes on the oily skin

behind your ear are more similar to those on your forehead — on the opposite side of your head — than to those on the dry skin of your forearm, just centimetres away.

Sebaceous sites — the forehead, the sides of the nose (the alar crease), the back, and the area behind the ears — are dominated by lipophilic (fat-loving) organisms, above all the genus *Cutibacterium* (formerly *Propionibacterium*). *Cutibacterium acnes*, the species most associated with acne vulgaris, thrives in sebaceous follicles — the tiny pits in the skin from which hairs emerge, surrounded by oil-producing sebaceous glands. These follicles create a sheltered, anaerobic, lipid-rich environment that *C. acnes* exploits with particular efficiency. It produces lipases that break down the triglycerides in sebum into free fatty acids, which it then uses as fuel. In this respect, *C. acnes* is to the sebaceous follicle what *Faecalibacterium prausnitzii* is to the colon: a specialist that has evolved to extract energy from a substrate that few competitors can match.

Think of sebaceous sites as oily ponds. The water (or in this case, the sebum) limits which species can flourish. The result is low diversity but high density — a few well-adapted species dominating overwhelmingly.

Moist sites — the armpit (axilla), the groin (inguinal crease), the inner elbow (antecubital fossa), the navel, and the spaces between the toes — are a different world. Here, temperature is higher, humidity is elevated, and the chemistry of the skin surface is shaped by sweat and the close apposition of skin folds. These conditions favour a broader range of bacteria, particularly *Staphylococcus* and *Corynebacterium* species. *Staphylococcus epidermidis*, the most abundant member of the genus on healthy skin, is especially well adapted to moist environments. *Corynebacterium* species are famous — or perhaps infamous — for their role in body odour: they metabolise the odourless precursors in apocrine sweat into the volatile compounds that give armpits their distinctive scent.

If sebaceous sites are oily ponds, moist sites are tropical lowlands — warm, humid, and supporting a richer diversity of life.

Dry sites — the forearm, the leg, the palm of the hand, the buttock — present the harshest conditions for microbial life on the skin. Exposed to air, low in moisture, subject to constant physical abrasion from clothing and contact, these are the deserts of the body surface. Microbial density is lower here than at sebaceous or moist sites, but taxonomic diversity is paradoxically higher. The Grice study found representatives from all four of the major skin-associated phyla — *Actinomycetota*, *Bacillota*, *Pseudomonadota*, and *Bacteroidota* — at dry sites, without any single group dominating as thoroughly as *Cutibacterium* dominates sebaceous areas [REF:grice2009].

The analogy holds: in a desert, no single species monopolises resources, because resources are scarce and widely distributed. Many species coexist, each scraping by on whatever is available.

Beyond these three broad categories, there are unique microbial micro-habitats scattered across the body. The **navel** (umbilicus) — a sheltered, poorly ventilated depression — harbours an unexpectedly diverse community that can include species rarely found elsewhere on the skin. The **toe web spaces** are warm, moist, and often occluded by footwear, creating conditions that favour not only bacteria but also fungi — including the dermatophytes responsible for athlete's foot. The **external ear canal** has its own characteristic community, shaped by the lipid-rich cerumen (earwax) that lines it.

The 2011 review by Grice and Segre, published in *Nature Reviews Microbiology*, synthesised these findings into a conceptual framework that has guided skin microbiome research ever since. The skin, they argued, should be understood not as a single organ but as a collection of distinct habitats — each with its own climate, its own chemistry, and its own characteristic community of microbial

inhabitants. The human body, in their analogy, is “a rich landscape of diverse environments that select for unique microbiomes” [REF:gricesegre2011].

This site-specificity has important clinical implications that we will return to throughout this chapter. The microbes involved in acne are not the same as those involved in eczema, and neither is the same as those involved in foot infections. Understanding skin disease requires understanding skin ecology — not just which microbe is present, but *where* it is, and what the local environment allows it to do.

6.2 Residents vs Transients — Who Belongs and Who’s Passing Through

Your skin is in constant contact with the outside world, which means it is constantly receiving microbial immigrants. Every surface you touch, every person you embrace, every pet you stroke, every door handle you grip deposits microbes on your skin. Yet the composition of your skin microbiome is remarkably stable over time. A study that re-sampled the same individuals months after their initial profiling found that the major features of their skin microbial communities were preserved — the same dominant species, in roughly the same proportions, at the same sites [REF:grice2009].

How is this possible, given the constant bombardment?

The answer lies in an ecological distinction as old as microbiology itself: the difference between **resident** and **transient** microorganisms.

Resident microbes are species that have established stable, self-sustaining populations on the skin. They are adapted to the local conditions — the pH, the moisture level, the available nutrients, the chemical defences of the host — and they persist for months or years, sometimes for an entire lifetime. When you wash your hands with soap and water, you remove the majority of transient organisms but only a fraction of the residents, which are anchored in deeper layers of the stratum corneum (the outermost layer of dead skin cells) and in the hair follicles beneath. Studies using progressively deeper skin sampling — from surface swabs to tape strips to punch biopsies — have shown that the deeper you sample, the more enriched the community becomes for resident species like *S. epidermidis* and *C. acnes*, which establish reservoirs in the follicular environment that surface washing cannot reach [REF:byrd2018].

Transient microbes are organisms deposited on the skin from the environment that do not establish permanent populations. They may survive on the surface for hours or days, but they lack the adaptations to compete with the established residents, and they are gradually displaced by the physical turnover of the skin itself — the constant shedding of dead cells from the stratum corneum, a process called desquamation, which removes an entire surface layer every two to three weeks.

The relationship between residents and transients is not always clear-cut. Some organisms are “occasional residents” — species that can persist for extended periods under certain conditions but do not maintain stable populations in all individuals. *Staphylococcus aureus*, for example, is carried persistently in the anterior nares (the front of the nostrils) by about 20 per cent of the population and intermittently by another 30 per cent. Whether it qualifies as a “resident” or a “transient” depends on the individual, the body site, and the time frame. In patients with atopic dermatitis, *S. aureus* can become a dominant resident of lesional skin — but this represents a disease-associated shift, not the healthy baseline.

The ecological concept that best describes how residents maintain their position is **colonisation**

resistance — the same principle we encountered in the gut (Chapter 5, section 5.6). An established, diverse community of resident microbes occupies the available niches, consumes the available nutrients, and in some cases actively produces substances that inhibit competitors. *Staphylococcus epidermidis*, for instance, produces antimicrobial peptides called **bacteriocins** that specifically target *S. aureus* and other potential pathogens. Several *Staphylococcus* species on healthy skin secrete **phenol-soluble modulins** and other small molecules that inhibit the growth of pathogenic organisms. The residents, in effect, patrol their territory.

This is not merely a metaphor. In a striking 2017 study, Teruaki Nakatsuji and colleagues at the University of California, San Diego, showed that coagulase-negative *Staphylococcus* strains isolated from the skin of healthy individuals — but not from patients with atopic dermatitis — produced antimicrobial peptides active against *S. aureus*. When these protective strains were transplanted back onto the skin of atopic dermatitis patients, *S. aureus* colonisation was significantly reduced [REF:nakatsuji2017]. The experiment demonstrated that the resident microbiome is not a passive bystander. It is an active participant in immune defence — and its composition can mean the difference between health and disease.

Where This Matters: The distinction between resident and transient organisms underpins the entire logic of hand hygiene in clinical settings. Handwashing with soap removes transients (including pathogens acquired from patients or contaminated surfaces) but leaves the resident flora largely intact. Alcohol-based hand sanitisers are more effective at killing transients but also do not eliminate residents sheltered in follicles and deeper skin layers. This is why healthcare-associated infections remain challenging despite rigorous hand hygiene protocols — and why understanding the skin's microbial ecology is directly relevant to hospital infection control.

6.3 The Skin Barrier and Microbial Defence

The skin faces a challenge analogous to the gut's but solved in a fundamentally different way. Both organs must prevent microbial invasion while coexisting with trillions of commensal organisms. The gut achieves this through a mucus layer — a soft, continuously secreted gel that separates microbes from the epithelial surface (as we saw in section 5.3). The skin has no mucus. Instead, it deploys a multi-layered defence system that is simultaneously physical, chemical, and immunological — and that works in partnership with the resident microbiota at every level.

The physical barrier is the stratum corneum — the outermost layer of the epidermis, consisting of 15 to 20 layers of dead, flattened keratinocytes (called corneocytes) embedded in a matrix of lipids, like bricks in mortar. This barrier is remarkably effective at preventing the entry of both microorganisms and water. Its integrity depends on the tight packing of the corneocytes and on the lipid matrix between them, which includes ceramides, cholesterol, and free fatty acids in a highly organised lamellar (layered) structure. Disruption of this lipid matrix — by harsh detergents, excessive washing, or genetic defects in the proteins that organise it — compromises the barrier and increases susceptibility to both infection and inflammation.

The most studied genetic defect affecting the skin barrier is loss-of-function mutation in the gene encoding **filaggrin** — a protein essential for the proper formation and hydration of the stratum corneum. Filaggrin mutations are present in approximately 10 per cent of the European population and are the strongest known genetic risk factor for atopic dermatitis (eczema). Individuals with

filaggrin mutations have a skin barrier that is more permeable, more prone to water loss, and more easily penetrated by allergens and microbes. The link between a structural defect in the barrier and a disease characterised by microbial dysbiosis (specifically, overgrowth of *S. aureus*) illustrates how intimately the physical barrier and the microbial ecosystem are connected [REF:byrd2018].

The chemical barrier includes the acidic pH of the skin surface (the so-called “acid mantle,” typically pH 4.5 to 5.5), which inhibits the growth of many pathogenic bacteria while favouring acid-tolerant residents like *S. epidermidis* and *Cutibacterium*. Free fatty acids in sebum — including sapienic acid and lauric acid — have direct antimicrobial activity. And the skin produces an arsenal of **antimicrobial peptides** (AMPs) that function as a kind of endogenous antibiotic system.

The most important skin AMPs include:

The **cathelicidin LL-37** — a broad-spectrum antimicrobial peptide produced by keratinocytes, sebocytes, and mast cells that is active against bacteria, fungi, and some viruses. LL-37 is upregulated in response to injury and infection, providing a rapid first-line chemical defence.

The **human beta-defensins** (hBD-1, hBD-2, hBD-3) — a family of small cysteine-rich peptides that kill microbes by disrupting their cell membranes. hBD-1 is constitutively expressed (always present), providing a baseline level of protection, while hBD-2 and hBD-3 are induced by microbial contact and inflammatory signals.

Dermcidin — an AMP secreted constitutively in sweat, providing antimicrobial defence across the entire skin surface. Unlike cathelicidin and the defensins, dermcidin does not require activation by infection or injury — it is part of the permanent chemical landscape of the skin.

Here is where the partnership between host and microbiota becomes explicit. The resident microbiota modulates AMP production. *S. epidermidis*, for example, has been shown to enhance keratinocyte production of AMPs, effectively amplifying the skin’s own chemical defences. At the same time, the resident microbes themselves produce antimicrobial compounds — the bacteriocins and phenol-soluble modulins mentioned in the previous section — that add another layer to the defence. The system is cooperative: the host provides the habitat and the baseline chemical environment; the residents contribute additional antimicrobial molecules that help exclude pathogens; and the pathogens must overcome both to establish infection.

The immunological barrier is the most complex layer. The skin contains a dense network of immune cells — Langerhans cells in the epidermis, dermal dendritic cells, resident T cells (including a large population of tissue-resident memory T cells), mast cells, macrophages, and innate lymphoid cells. These cells continuously sample the microbial landscape, discriminating between familiar commensals and unfamiliar invaders.

The immune system’s relationship with the skin microbiota is not merely tolerant. It is actively shaped by microbial contact. Studies in germ-free mice — animals raised without any microorganisms — show that the skin immune system develops abnormally in the absence of commensals. Germ-free mice have fewer skin-resident T cells, produce fewer AMPs, and mount weaker inflammatory responses to skin pathogens. Colonisation with specific commensal species, particularly *S. epidermidis*, restores many of these deficits. This is immune education — the same principle we will explore in greater depth in Chapter 12 — and it demonstrates that the skin’s defences are not fully functional without their microbial partners [REF:byrd2018].

The three layers — physical, chemical, and immunological — do not operate in isolation. They form an integrated defence system in which the microbiota is both protected by the barrier and essential

to its function. Damage the physical barrier (as in filaggrin deficiency), and the microbial community shifts toward pathogenic species. Suppress the chemical defences (as happens with prolonged corticosteroid use, which reduces AMP production), and opportunistic infections become more likely. Disrupt the immune education provided by commensals (as in germ-free conditions), and the entire system weakens.

The skin barrier, in other words, is not just a wall. It is an ecosystem — and the microbes that live on it are structural components of that wall, not invaders trying to breach it.

6.4 *Malassezia* and the Skin Mycobiome — The Dominant Fungal Genus Hiding in Plain Sight

When people think of the skin microbiome, they typically think of bacteria. This is understandable — bacteria are the most abundant and the most studied organisms on the skin surface. But the skin, like the gut, is home to a multi-kingdom community that includes fungi, viruses, and mites. Among these, the fungi deserve particular attention, because the fungal ecology of the skin turns out to be both simpler and stranger than anyone expected.

The study that changed our understanding of the skin mycobiome was published in *Nature* in 2013 by Keisha Findley, Julia Segre, and colleagues — the same NIH group that had mapped the bacterial geography of the skin four years earlier. Using fungal-specific DNA sequencing (targeting the internal transcribed spacer, or ITS region, of fungal ribosomal RNA genes), they surveyed 14 skin sites across 10 healthy adults. What they found was striking in its simplicity: 11 of the 14 sites — essentially all of the core body and arm sites — were overwhelmingly dominated by a single fungal genus: *Malassezia* [REF:findley2013].

Malassezia is a genus of basidiomycete yeasts that are uniquely adapted to life on human skin. They are **lipid-dependent** — unlike most fungi, they cannot synthesise their own fatty acids and must obtain them from their environment. On human skin, that environment is sebum: the oily secretion of the sebaceous glands. *Malassezia* species possess lipase enzymes that break down the triglycerides in sebum, liberating the fatty acids they need for growth. This metabolic dependence explains why *Malassezia* is most abundant at sebaceous sites (the scalp, face, chest, and back) and relatively sparse at dry sites like the forearm.

The dominance of *Malassezia* on the skin is remarkable for several reasons. First, it represents one of the most extreme examples of single-genus dominance in any human-associated microbial community. While the gut harbours hundreds of bacterial genera, the skin mycobiome at most body sites is essentially a *Malassezia* monoculture, with other genera present only at trace levels. Second, this dominance is consistent across individuals, across ethnicities, and across geographic regions — suggesting a very deep evolutionary association between *Malassezia* and human skin.

The exception to this *Malassezia* monopoly was equally revealing. The three foot sites surveyed — plantar heel, toenail, and toe web space — showed a dramatically different picture: high fungal diversity, with genera including *Aspergillus*, *Cryptococcus*, *Rhodotorula*, *Epicoccum*, and *Cladosporium* alongside *Malassezia*. The feet, in other words, are a fungal garden, while the rest of the body is a *Malassezia* plantation. This pattern likely reflects the foot's unique environment: enclosed in footwear, warm and moist, in constant contact with environmental surfaces (floors, soil, shower stalls) that harbour a diverse fungal community [REF:findley2013].

Malassezia occupies an ambiguous position in human health. In most people, most of the time, it is a commensal — a harmless or even beneficial member of the skin community. But several *Malassezia* species are associated with common skin conditions. *Malassezia globosa* and *M. restricta* are the primary agents of **dandruff** and **seborrheic dermatitis** — the flaky, sometimes itchy scaling of the scalp and other sebaceous areas that affects up to 50 per cent of adults. The mechanism involves *Malassezia*'s lipase activity: the free fatty acids released by the breakdown of sebum triglycerides — particularly oleic acid — irritate the skin of susceptible individuals, triggering an inflammatory response that accelerates the turnover of epidermal cells and produces visible flaking. This is why antifungal shampoos containing ketoconazole or zinc pyrithione — which reduce *Malassezia* populations — are effective treatments for dandruff.

Malassezia is also implicated in **pityriasis versicolor** (tinea versicolor) — a common condition in which patches of skin become lighter or darker than the surrounding area, caused by *Malassezia* species growing in their hyphal (filamentous) form in the stratum corneum. The hyphal form is associated with disease, while the yeast form is the commensal state — a pattern reminiscent of *Candida albicans* in the gut, where the yeast-to-hypha transition is similarly associated with pathogenicity.

More recently, *Malassezia* has attracted attention for its potential role in **atopic dermatitis**. Several studies have found elevated *Malassezia*-specific IgE antibodies in patients with atopic dermatitis, particularly those with head and neck involvement. The hypothesis is that in individuals with a compromised skin barrier (for example, due to filaggrin mutations), *Malassezia* antigens penetrate deeper into the skin, triggering an allergic immune response that exacerbates the eczema. This remains an active area of research, and the relationship between *Malassezia* and atopic dermatitis is not yet fully established — but it illustrates the broader principle that organisms which are harmless in one context can contribute to disease when the barrier breaks down.

The discovery that the skin mycobiome is dominated by a single genus raises a question that mycologists are still working to answer: why? One hypothesis is that *Malassezia*'s unique lipid dependence gives it an insurmountable competitive advantage in the lipid-rich environment of sebaceous skin. Another is that the human immune system, through long co-evolution, has developed a tolerance to *Malassezia* that it does not extend to other fungal genera — effectively creating an immunological niche that only *Malassezia* can occupy. A third possibility is that *Malassezia* actively suppresses the growth of competing fungi through the production of antifungal metabolites. The truth is probably some combination of all three.

Where This Matters: If you have ever used an antifungal shampoo for dandruff, you have conducted an experiment on your skin mycobiome. The treatment works not by sterilising the scalp but by reducing *Malassezia* numbers enough to lower the concentration of irritant fatty acids below the threshold that triggers inflammation. When you stop treatment, *Malassezia* regrows — because it is a permanent resident — and dandruff often returns. This is mycobiome management, not mycobiome elimination.

6.5 Acne, Eczema, and Wound Healing — When the Ecosystem Breaks Down

If the previous sections described the skin microbiome in health, this one examines what happens when the balance tilts — when the carefully maintained partnership between host and mi-

crobes breaks down, and the ecological community shifts from a state of resilience to one of disease. Three conditions illustrate this breakdown with particular clarity: acne vulgaris, atopic dermatitis (eczema), and chronic wounds.

Acne Vulgaris: Not Simply an Infection

Acne is the most common skin condition on Earth, affecting an estimated 85 per cent of adolescents and persisting well into adulthood in many individuals. For decades, it was framed as an infection — a simple story in which *Propionibacterium acnes* (now reclassified as *Cutibacterium acnes*) invades the sebaceous follicle and causes inflammation. This narrative justified the widespread use of antibiotics as a first-line treatment for acne, a practice that continues today.

The microbiome era has complicated this story considerably.

The problem with the “infection” model is that *C. acnes* is present on the skin of virtually every post-pubertal human, whether or not they have acne. It is a universal resident of sebaceous follicles, not an invading pathogen. If *C. acnes* caused acne simply by being present, everyone would have it — and while acne is common, it is not universal.

The key insight from genomic studies is that *C. acnes* is not one organism but many. Strain-level analysis — examining the genetic differences between individual *C. acnes* isolates — has revealed that specific **ribotypes** (genetic subtypes) of *C. acnes* are enriched in acne lesions, while other ribotypes are more common on healthy skin. The disease-associated strains tend to produce more inflammatory molecules, including a porphyrin compound that generates reactive oxygen species and triggers an immune response. In contrast, healthy-skin strains produce fewer of these inflammatory mediators and may even have anti-inflammatory properties [REF:byrd2018].

Acne, in this newer view, is not an infection but an ecological shift — a change in the *strain-level* composition of *C. acnes* within the follicle, combined with hormonal changes (particularly the rise in androgen-driven sebum production during puberty) and host immune responses. The condition arises at the intersection of microbial ecology, endocrinology, and immunology — not from any one factor alone.

This has practical implications. The longstanding practice of treating acne with broad-spectrum antibiotics — tetracyclines, macrolides — may reduce *C. acnes* overall, but it does not discriminate between disease-associated and health-associated strains, and it carries the well-known risks of promoting antibiotic resistance and disrupting the broader skin and gut microbiomes. Emerging strategies, still largely experimental, aim to be more precise: selectively targeting the virulent strains while preserving the commensals, or even transplanting health-associated *C. acnes* strains to shift the follicular community back toward a non-inflammatory state.

Atopic Dermatitis: A Microbial Signature of Flare

Atopic dermatitis (AD) — commonly called eczema — is a chronic inflammatory skin disease characterised by intense itching, dry skin, and recurrent flares of red, weeping, or crusted lesions. It affects up to 20 per cent of children and 3 per cent of adults in industrialised countries, and its prevalence has increased dramatically over the past half century — a trend that parallels the rise of other allergic and autoimmune conditions and that we will explore in section 6.6.

The microbial signature of atopic dermatitis is more dramatic than that of acne. In a landmark 2012 study, Heidi Kong and colleagues at the NIH tracked the skin microbiome of children with AD

through cycles of flare and remission. What they found was a pattern of striking simplicity: during disease flares, the microbial diversity of affected skin collapsed, and *Staphylococcus aureus* surged to dominance — in some cases accounting for more than 90 per cent of all bacteria on lesional skin. When the flare resolved, diversity recovered, and *S. aureus* abundance declined [REF:kong2012].

This is dysbiosis in its purest form: a shift from a diverse, balanced community to one dominated by a single, potentially pathogenic species. *S. aureus* is not merely a bystander in atopic dermatitis. It actively worsens the disease through several mechanisms. It produces **superantigens** — toxins that non-specifically activate large numbers of T cells, amplifying the inflammatory response. It secretes **proteases** that damage the already compromised skin barrier. It forms **biofilms** that resist both the immune system and topical treatments. And its presence correlates with disease severity: the more *S. aureus*, the worse the eczema.

But *S. aureus* overgrowth is not just a consequence of the barrier defect. It is also a consequence of losing the commensal bacteria that normally keep it in check. Remember the colonisation resistance we discussed in section 6.2: on healthy skin, *S. epidermidis* and other coagulase-negative staphylococci produce antimicrobial molecules that inhibit *S. aureus*. In atopic dermatitis, these protective commensals are depleted during flares, removing a critical layer of ecological defence [REF:kong2012].

This understanding has opened a genuinely new therapeutic frontier. If the problem is not just *S. aureus* overgrowth but the loss of protective commensals, then the treatment might not be antibiotics (which kill commensals too) but rather the **restoration of the missing microbial community**. Early clinical trials of topical application of commensal *Staphylococcus* species — a form of “bacteriotherapy” — have shown promising results, reducing *S. aureus* colonisation and improving disease severity in patients with atopic dermatitis. This approach — treating a skin disease by adding microbes rather than killing them — would have seemed absurd twenty years ago. It is now one of the most active areas of translational skin microbiome research.

Chronic Wounds: When Healing Stalls

Acute wounds — cuts, scrapes, surgical incisions — normally heal through a well-orchestrated series of phases: haemostasis (clotting), inflammation, proliferation (new tissue growth), and remodelling. Microbes are present throughout this process, and in most cases, the resident skin flora re-colonises the healing wound without causing infection.

Chronic wounds are different. These are wounds that fail to progress through the normal healing sequence and remain open for weeks, months, or even years. They include diabetic foot ulcers, venous leg ulcers, and pressure ulcers (bedsores), and they affect an estimated 6.5 million people in the United States alone. They are enormously costly — both financially and in terms of suffering.

The microbiome of chronic wounds is characterised by the presence of **biofilms** — structured, multi-species microbial communities encased in a self-produced matrix of sugars and proteins that adheres to the wound surface and is extraordinarily resistant to both antibiotics and the host immune response. Biofilms are found in an estimated 60 to 90 per cent of chronic wounds but only 6 per cent of acute wounds — a disparity that strongly implicates biofilm formation in the failure to heal.

The species composition of chronic wound biofilms varies, but common members include *S. aureus*, *Pseudomonas aeruginosa*, anaerobic species like *Peptostreptococcus* and *Finegoldia*, and increasingly, fungi. These are not random infections. They are organised communities with division of

labour, chemical communication (quorum sensing), and resistance mechanisms that make them fundamentally different from the same species growing as free-floating (planktonic) cells.

The relationship between the wound microbiome and healing failure is not straightforward. Not all colonised wounds become infected, and not all biofilm-containing wounds fail to heal. The current view is that the wound microbiome exists on a spectrum: from benign colonisation (commensal organisms present but not interfering with healing) through critical colonisation (microbial load increasing, healing slowing) to overt infection (invasion of viable tissue, systemic inflammatory response). Where a particular wound sits on this spectrum depends on the balance between microbial virulence and host defences — and in patients with diabetes, vascular disease, or immunosuppression, the balance is tipped in the microbes' favour.

Understanding chronic wound biofilms has changed clinical practice. Physical disruption of the biofilm — through debridement (the removal of dead tissue and biofilm material from the wound surface) — is now considered essential for chronic wound management, because antibiotics alone cannot penetrate the biofilm matrix effectively. The principle is ecological: you cannot treat the infection without disrupting the community structure.

6.6 The Hygiene Hypothesis Revisited

We close this chapter with a question that connects the skin microbiome to one of the most debated and frequently misunderstood ideas in modern medicine — and one that extends far beyond dermatology.

In 1989, the epidemiologist David Strachan published a short paper in the *British Medical Journal* reporting a seemingly trivial observation: among a cohort of 17,414 British children followed since birth, the risk of developing hay fever was inversely associated with the number of older siblings in the household. Children with more older brothers and sisters were *less* likely to develop allergic rhinitis. Strachan proposed, in a single speculative paragraph, that “declining family size, improvements in household amenities, and higher standards of personal cleanliness have reduced the opportunity for cross infection in young families” and that this reduction in early childhood infections might explain the rising prevalence of allergic disease [REF:strachan1989].

This was the birth of the **hygiene hypothesis** — and the name has been causing confusion ever since.

The confusion arises because the name implies that the solution to allergic disease is less hygiene — less handwashing, less cleanliness, less sanitation. This is not what the evidence supports. The decades of research that followed Strachan's observation have progressively refined the hypothesis, and the modern version looks quite different from the original.

The key refinement came from immunologist Graham Rook, who in the early 2000s proposed what he called the “**Old Friends**” hypothesis. Rook argued that the relevant missing exposures were not childhood infections (measles, mumps, and the like) but rather the organisms with which humans co-evolved over millions of years — environmental mycobacteria, helminths (parasitic worms), and the diverse microbiota encountered through close contact with soil, animals, and unprocessed food. These “old friends,” Rook argued, played a crucial role in calibrating the immune system during development — specifically, in training the **regulatory T cells** that suppress inappropriate immune responses. Without sufficient early exposure to these immunomodulatory or-

ganisms, the immune system develops a bias toward overreaction — manifesting as allergy, asthma, and autoimmune disease [REF:rook2014].

The Finnish ecologist Ilkka Hanski and immunologist Tari Haahtela extended this further with the **biodiversity hypothesis**, published in 2012. Drawing on data from a large study of adolescents in eastern Finland, they showed that individuals living in environments with greater plant biodiversity — surrounded by more species of flowers and native vegetation — had more diverse skin microbiomes and lower rates of atopic sensitisation (allergic immune responses). The mechanism, they proposed, was environmental: biodiverse natural surroundings expose the skin to a richer community of environmental microorganisms, which in turn promotes immune regulation. Urbanisation, by replacing diverse natural environments with concrete and monoculture lawns, reduces this microbial exposure — and the immune system suffers the consequences.

What does this mean for the skin specifically?

The skin is the primary interface between the body and the external microbial environment. It is the surface through which much of the “Old Friends” exposure occurs — through contact with soil, water, animals, other people, and the diverse microbial communities that these carry. If the biodiversity hypothesis is correct, then the skin microbiome is not just a reflection of personal hygiene. It is a reflection of the **ecological richness of the environment** in which a person lives.

This has been supported by several lines of evidence. Children raised on farms — in close contact with livestock, hay, and unpasteurised milk — have lower rates of asthma and allergic disease than their urban peers, a finding consistently replicated across European studies. These farm-raised children also have more diverse gut and skin microbiomes, and their immune cells show stronger regulatory activity. The effect appears to require early exposure: moving to a farm as an adult does not confer the same protection as growing up on one.

The practical implications are complex. The hygiene hypothesis, properly understood, does *not* argue against handwashing, vaccination, food safety, or sanitation — all of which prevent genuine infections that cause suffering and death. What it argues is that modern urban life, by reducing contact with environmental microbial diversity, may deprive the developing immune system of the regulatory signals it needs to function properly. The solution is not to abandon hygiene but to find ways to restore the microbial exposures that hygiene need not eliminate — through contact with nature, diverse diets, reduced unnecessary antibiotic use, and perhaps, eventually, through rationally designed microbial interventions.

The skin, as the body’s outer boundary, sits at the centre of this story. It is where the outside world meets the immune system. It is where microbial diversity — or its absence — has its most immediate impact. And it is where the consequences of our increasingly sanitised, urbanised, biodiversity-depleted modern environment are written, quite literally, on the body’s surface.

Where This Matters: The next time someone tells you that you should let your children play in the dirt because “it’s good for their immune system,” they are drawing — however loosely — on three decades of research into the hygiene hypothesis, the Old Friends mechanism, and the biodiversity hypothesis. The underlying science is far more nuanced than the folk wisdom suggests, but the core intuition is supported by evidence: early, diverse microbial exposure through the skin and gut helps calibrate the immune system. The question is no longer whether this matters, but how to restore it in a world that has been built to minimise it.

In the next chapter, we move from the body's outer surface to its inner airways — the respiratory tract, long assumed to be sterile, now known to harbour its own complex microbial community. The lung microbiome, like the skin's, is shaped by its local environment, disrupted in disease, and far more important to health than anyone suspected.

Chapter 6 — Key References

- [REF:grice2009] Grice, E.A., Kong, H.H., Conlan, S., Deming, C.B., Davis, J., Young, A.C., ... & Segre, J.A. (2009). Topographical and temporal diversity of the human skin microbiome. *Science*, 324(5931), 1190–1192. DOI: 10.1126/science.1171700
- [REF:gricesegre2011] Grice, E.A. & Segre, J.A. (2011). The skin microbiome. *Nature Reviews Microbiology*, 9(4), 244–253. DOI: 10.1038/nrmicro2537
- [REF:byrd2018] Byrd, A.L., Belkaid, Y. & Segre, J.A. (2018). The human skin microbiome. *Nature Reviews Microbiology*, 16(3), 143–155. DOI: 10.1038/nrmicro.2017.157
- [REF:findley2013] Findley, K., Oh, J., Yang, J., Conlan, S., Deming, C., Meyer, J.A., ... & Segre, J.A. (2013). Topographic diversity of fungal and bacterial communities in human skin. *Nature*, 498(7454), 367–370. DOI: 10.1038/nature12171
- [REF:kong2012] Kong, H.H., Oh, J., Deming, C., Conlan, S., Grice, E.A., Beatson, M.A., ... & Segre, J.A. (2012). Temporal shifts in the skin microbiome associated with disease flares and treatment in children with atopic dermatitis. *Genome Research*, 22(5), 850–859. DOI: 10.1101/gr.131029.111
- [REF:nakatsuji2017] Nakatsuji, T., Chen, T.H., Narala, S., Chun, K.A., Two, A.M., Yun, T., ... & Gallo, R.L. (2017). Antimicrobials from human skin commensal bacteria protect against *Staphylococcus aureus* and are deficient in atopic dermatitis. *Science Translational Medicine*, 9(378), eaah4680. DOI: 10.1126/scitranslmed.aah4680
- [REF:strachan1989] Strachan, D.P. (1989). Hay fever, hygiene, and household size. *BMJ*, 299(6710), 1259–1260. DOI: 10.1136/bmj.299.6710.1259
- [REF:rook2014] Rook, G.A.W., Raison, C.L. & Lowry, C.A. (2014). Microbiota, immunoregulatory old friends and psychiatric disorders. *Advances in Experimental Medicine and Biology*, 817, 319–356. DOI: 10.1007/978-1-4939-0897-4_15

Additional references to be verified and expanded during manuscript preparation.

Chapter 7: The Respiratory Tract — Not So Sterile After All

In the previous chapter, we explored the skin — the body's outermost surface, constantly exposed to the environment yet home to stable, site-specific microbial communities. Now we turn inward, to a system that is at once more sheltered and more vulnerable: the respiratory tract.

Every breath you take draws in roughly half a litre of air, and with it, several hundred to several thousand microorganisms — bacteria, fungal spores, and viruses carried on dust particles, aerosol droplets, and the invisible biological haze that surrounds every living thing. Over the course of a day, your lungs process something like 10,000 litres of air. The respiratory tract is not a sealed chamber. It is an open system, in continuous dialogue with the microbial world.

And yet, for most of the history of modern medicine, the lungs were considered sterile. The lower respiratory tract — everything below the vocal cords — was thought to be a pristine, microbe-free zone in healthy people, defended by a gauntlet of physical and immunological barriers so effective that no organism could survive there under normal conditions. This belief was stated as fact in textbooks for the better part of a century, and it shaped clinical thinking in ways that are still being untangled today.

It was wrong.

The lungs are not sterile. They harbour a sparse but real microbial community — a **lung microbiome** — that is dynamically maintained by a balance between microbial immigration and elimination. This community is far less dense than the gut's, far less stable than the skin's, and far more recently discovered than either. But its existence has fundamentally changed how we think about respiratory health and disease.

This chapter follows the respiratory tract from its outermost gateway — the nose and nasopharynx — down through the airways to the alveoli deep in the lungs. We will explore how the “sterile lung” dogma fell, what the healthy lung microbiome actually looks like, how the nasopharynx serves as both gatekeeper and reservoir, and what happens to the respiratory microbiome in three of the most important chronic lung diseases: chronic obstructive pulmonary disease (COPD), asthma, and cystic fibrosis. We will then meet the fungi of the respiratory tract — *Aspergillus* in particular — and close with the lessons of the COVID-19 pandemic, which provided the most dramatic demonstration in living memory of what happens when a viral infection shatters the delicate microbial ecology of the lungs.

7.1 Overturning the Sterile Lung Dogma

The idea that healthy lungs are sterile was never really proven. It was assumed.

The assumption made a certain kind of sense. The respiratory tract is equipped with a formidable series of defences designed to prevent microbial colonisation. Air entering the nose is filtered by nasal hairs and warmed and humidified by the turbinates — the scrolled bony structures of the nasal passages. Particles larger than about ten micrometres are trapped in the nasal mucus. Smaller

particles that penetrate further are caught by the **mucociliary escalator** — a continuous carpet of mucus propelled upward by the synchronised beating of cilia (hair-like projections) on the epithelial cells lining the airways. This escalator operates around the clock, sweeping trapped particles and microbes up from the lower airways toward the throat, where they are swallowed into the acid bath of the stomach. Below this, the airways narrow progressively, and turbulent airflow gives way to smooth laminar flow in the smallest bronchioles, further reducing the chance that inhaled particles will deposit on the airway walls. Finally, any microbes that do reach the alveoli — the tiny air sacs where gas exchange occurs — are met by **alveolar macrophages**, specialised immune cells that patrol the alveolar surface and engulf and destroy microbial intruders.

With all of these defences in place, it seemed reasonable to conclude that healthy lungs would be free of bacteria. And when early microbiologists attempted to culture bacteria from the lower respiratory tract, they usually found nothing — or, when they did find something, they attributed it to contamination from the mouth and throat during sampling. This was a legitimate concern: the only way to sample the lower airways is to pass a bronchoscope through the mouth or nose, traversing the heavily colonised upper airways in the process. Any bacteria detected could plausibly have been carried down from above.

The result was a kind of epistemic stalemate. Positive cultures were dismissed as contamination, and negative cultures were accepted as confirmation of sterility. The conclusion — sterile lungs — reinforced itself. And because the assumption of sterility seemed so well established, the lungs were not included when the National Institutes of Health launched the **Human Microbiome Project** (HMP) in 2007. The HMP surveyed the microbiome of the gut, the skin, the mouth, the nose, and the urogenital tract. The lower respiratory tract was excluded — not because anyone had rigorously demonstrated its sterility, but because it was simply taken as given [REF:mammen2016].

The dogma began to crumble with the arrival of **culture-independent sequencing**. As we discussed in Chapter 4, the ability to identify microbes by their DNA — rather than by growing them in a laboratory dish — transformed microbiology. For the lungs, this was decisive. Culture-based methods require microbes that can grow under laboratory conditions, and many of the organisms found in the lungs are slow-growing, oxygen-sensitive, or otherwise difficult to cultivate. DNA-based methods detect them regardless.

The study widely credited with breaking the sterile lung paradigm was published in 2010 by Markus Hilty, William Cookson, Miriam Moffatt, and colleagues at Imperial College London. Using 16S ribosomal RNA gene sequencing — the same molecular barcode technique that had revealed the complexity of the gut and skin microbiomes — they analysed bronchoalveolar lavage (BAL) samples from the lower airways of adults with asthma, adults with COPD, and healthy controls. The finding that changed the field was this: the bronchial tree of healthy individuals was not sterile. It contained a detectable microbial community — sparse compared to the mouth or the gut, but real, diverse, and structured. The researchers identified approximately 2,000 bacterial genomes per square centimetre of airway surface. The composition differed between health and disease: healthy lungs were enriched in members of the phylum **Bacteroidota** (particularly *Prevotella*), while the lungs of asthmatic and COPD patients were enriched in **Proteobacteria** (particularly *Haemophilus*) [REF:hilty2010].

This was a landmark moment. For the first time, using rigorous culture-independent methods with controls for contamination, someone had demonstrated that the healthy human lung has a microbiome.

The finding was not without controversy. Critics — and there were many — raised the contamina-

tion question again: could the 16S sequences detected in BAL fluid have come from the bronchoscope's passage through the upper airways, or from reagent contamination during DNA extraction and amplification (the so-called “kitome” problem)? These were not trivial objections, and addressing them required years of methodological refinement.

The most important conceptual advance came from Robert Dickson, Gary Huffnagle, and colleagues at the University of Michigan, who reframed the question entirely. In a series of influential papers between 2013 and 2016, Dickson proposed the “**adapted island model**” of lung biogeography. The key insight was ecological: the lung microbiome should not be thought of as a fixed, resident community in the way the gut microbiome is. Instead, it is a dynamic equilibrium — a balance between three processes: **immigration** of microbes into the lungs (primarily by microaspiration from the mouth and throat, and by inhalation of airborne organisms), **elimination** of microbes from the lungs (by mucociliary clearance, coughing, and immune killing), and **local reproduction** of those organisms that can grow in the lung environment [REF:dickson2015a].

In health, immigration and elimination are roughly balanced, and local reproduction is minimal. The lung microbiome is therefore sparse, transient, and dominated by organisms from the upper airways — particularly from the mouth. It is, in Dickson's analogy, like an island whose species composition is determined mainly by immigration from a nearby mainland (the oral cavity) rather than by evolution of species on the island itself. The community turns over continuously; the same species persist not because they are stably established but because they are constantly being replenished from above [REF:dickson2015b].

In disease, this balance tips. Inflammation alters the lung environment — increasing temperature, providing new nutrients from damaged tissue, reducing oxygen tension in congested areas, and impairing the clearance mechanisms. These changes favour the growth of specific organisms, allowing them to reproduce locally and shift the community from one that is immigration-dominated to one that is growth-dominated. This is when the sparse, diverse background community gives way to the pathogen-dominated, low-diversity communities seen in pneumonia, COPD exacerbations, and other lung diseases.

The adapted island model resolved many of the earlier controversies. It explained why the lung microbiome resembles the oral microbiome (because most of its members arrive from the mouth via microaspiration). It explained why the community is so sparse compared to the gut (because elimination is rapid and local growth conditions are harsh). And it explained why disease states look so different (because disease tips the balance toward local reproduction of organisms that thrive in the altered environment).

By the mid-2010s, the sterile lung dogma was effectively dead — replaced by an ecological understanding of the respiratory tract as a dynamic microbial ecosystem, far sparser than the gut but no less interesting and no less important to health and disease.

Where This Matters: If you have ever had a bronchoscopy — or watched one performed — you will appreciate the challenge of sampling the lower airways without contamination. The bronchoscope passes through the mouth and throat, where billions of bacteria coat every surface. The development of protected specimen brushes, careful BAL protocols, and rigorous sequencing controls was essential to distinguishing genuine lung microbiome signals from upper-airway carry-over. The fact that the lung microbiome was hidden in plain sight for decades is a reminder that what we fail to find is sometimes a failure of method, not of nature.

7.2 The Nasopharynx as a Microbial Gateway

If the lung microbiome is an island, the nasopharynx is its mainland.

The **nasopharynx** — the space behind the nose and above the soft palate — occupies an anatomically unique position in the respiratory system. It sits at the junction of the nasal passages, the oral cavity, the middle ears (via the Eustachian tubes), and the lower airways. Every breath you draw passes through it. Every microbe that enters your respiratory tract encounters it. It is, in a very literal sense, the gateway between the outside world and the lungs.

And it is far from sterile. The nasopharynx is densely colonised, beginning within hours of birth, and it harbours a microbial community whose composition has profound implications for respiratory health — not just locally, but throughout the airways.

The nasopharyngeal microbiome has been characterised in detail over the past decade, thanks to large cohort studies using 16S rRNA sequencing. What these studies reveal is a community organised around a small number of dominant bacterial genera, each of which characterises a distinct **microbiome profile** (sometimes called a “biotype” or “cluster”). The most commonly identified profiles are dominated by one of the following genera: *Moraxella*, *Streptococcus*, *Haemophilus*, *Corynebacterium*, *Dolosigranulum*, or *Staphylococcus* [REF:cleary2017].

These profiles are not random. They show consistent associations with health and disease.

Profiles dominated by **Corynebacterium** and **Dolosigranulum** — both regarded as commensal organisms — are consistently associated with respiratory health. Children whose nasopharynx is dominated by these genera have lower rates of acute respiratory infections, including otitis media (middle-ear infections), bronchiolitis, and pneumonia. The mechanism is thought to involve both direct competition with pathogens and modulation of local immune responses. *Dolosigranulum pigrum*, in particular, has attracted attention as a potentially protective species — a commensal that may actively suppress the growth of pathogenic neighbours.

By contrast, profiles dominated by **Haemophilus** (principally *H. influenzae*), **Streptococcus** (principally *S. pneumoniae*), or **Moraxella** (principally *M. catarrhalis*) are associated with increased susceptibility to respiratory infections. These three genera — sometimes called the “pathobiont trio” — are familiar to any clinician who treats respiratory disease. They are the most common bacterial causes of otitis media, sinusitis, and community-acquired pneumonia. But here is the key insight from microbiome studies: these organisms are not simply absent in health and present in disease. They are present in the nasopharynx of many healthy individuals, held in check by the surrounding community. Disease occurs not because the pathogen arrives but because the ecological balance shifts — through viral infection, antibiotic use, or other disruptions — allowing the pathobiont to proliferate and, in some cases, to descend into the lower airways.

A landmark study by Wouter de Steenhuijsen Piters, Debby Bogaert, and colleagues in 2016 demonstrated this interaction directly. They profiled the nasopharyngeal microbiome of young children with RSV (respiratory syncytial virus) infections of varying severity and compared them with healthy controls. Five distinct microbiome clusters emerged, dominated respectively by *Haemophilus influenzae*, *Streptococcus*, *Corynebacterium*, *Moraxella*, and *Staphylococcus aureus*. RSV infection and hospitalisation were positively associated with *Haemophilus*- and *Streptococcus*-dominated profiles, and negatively associated with *Staphylococcus aureus* abund-

ance. Crucially, the microbiome profile did not merely correlate with disease severity — it was associated with distinct host immune responses. Children with RSV infections and *Haemophilus*- or *Streptococcus*-dominated nasopharyngeal communities showed greater overexpression of genes linked to Toll-like receptor signalling, neutrophil activation, and macrophage signalling, suggesting that the nasopharyngeal microbiome modulates the host's inflammatory response to viral infection [REF:desteenhuijsen2016].

This finding has a practical implication that deserves emphasis. Much of the morbidity and mortality from respiratory viral infections — RSV in infants, influenza in the elderly — comes not from the virus itself but from secondary bacterial infections. The composition of the nasopharyngeal microbiome may determine which individuals are at highest risk of these bacterial complications. A child whose nasopharynx is dominated by *Corynebacterium* and *Dolosigranulum* may weather an RSV infection with nothing more than a runny nose. A child whose nasopharynx is dominated by *Haemophilus* or *Streptococcus* may develop otitis media, bronchiolitis, or pneumonia from the same viral exposure. The virus opens the door; the pre-existing microbial community determines what walks through it.

The nasopharynx also acts as a **reservoir** — a standing pool of organisms that can seed the lower airways. As we noted in the previous section, microaspiration from the upper airways is the primary route by which microbes enter the lungs. Most of this aspiration occurs during sleep, when the swallowing reflex is suppressed and small volumes of secretions from the nasopharynx and oropharynx trickle past the vocal cords. In healthy individuals, this is inconsequential — the aspirated organisms are quickly cleared. But in disease — particularly in conditions that impair the mucociliary escalator, reduce cough effectiveness, or compromise immune function — microaspiration becomes a highway for pathogen delivery to the lungs. Mechanical ventilation in the intensive care unit, for instance, bypasses the body's natural airway defences entirely, allowing oropharyngeal organisms direct access to the lower airways. This is why ventilator-associated pneumonia is one of the most common and dangerous hospital-acquired infections.

The relationship between the nasopharynx and the lungs is therefore not one-directional. The nasopharynx shapes the lung microbiome by providing its source organisms. The lung's local conditions — inflammation, mucus production, oxygen levels — then select which of those organisms thrive or perish. And the composition of the resulting lung community, in turn, influences the host immune response, potentially amplifying or suppressing the very inflammation that determines whether the community remains benign or becomes pathogenic.

Understanding this gateway function has clinical implications that extend beyond respiratory medicine. Nasal and nasopharyngeal swabs — cheap, non-invasive, and easily obtained — may serve as accessible windows into respiratory health and disease risk. If the nasopharyngeal microbiome predicts susceptibility to secondary bacterial infection during viral illness, as the evidence increasingly suggests, then profiling it could inform clinical decision-making in ways that current diagnostic tools cannot.

7.3 The Lung Microbiome in Health and Disease

In a healthy lung, the microbiome is sparse, diverse, and dominated by organisms carried in from the mouth and nasopharynx. The most commonly detected genera in bronchoalveolar lavage from healthy individuals are **Prevotella**, **Veillonella**, **Streptococcus**, and **Fusobacterium** — all of

which are abundant in the oral cavity. Their presence in the lungs reflects the constant low-level microaspiration that everyone experiences, particularly during sleep. The total bacterial burden in healthy lungs is orders of magnitude lower than in the gut — roughly 10^3 to 10^5 organisms per millilitre of BAL fluid, compared to 10^{11} per gram of colonic content [REF:dickson2016].

But this sparse community is not inert. In a striking series of studies, Leopoldo Segal and colleagues at New York University showed that some healthy individuals harbour a lung microbiome enriched in oral taxa — particularly *Prevotella* and *Veillonella* — that they designated **pneumotypeSPT** (for supraglottic-predominant taxa). Compared with individuals whose BAL was essentially sterile (resembling the saline background controls), those with pneumotypeSPT had significantly higher numbers of lymphocytes and neutrophils in their BAL fluid, elevated levels of inflammatory cytokines, and a distinct pro-inflammatory immune profile characterised by **Th17** lymphocyte activation. In other words, even in the absence of any clinical disease, the presence of oral microbes in the lungs was associated with measurable subclinical inflammation [REF:segal2016].

This observation is important for several reasons. It suggests that the lung microbiome, sparse as it is, is not immunologically invisible. The immune system “sees” the microbes that reach the alveoli and responds to them. In most people, this response is appropriately calibrated — sufficient to clear the aspirated organisms without causing tissue damage. But in individuals with impaired clearance or enhanced aspiration (smokers, the elderly, patients with neurological conditions that impair swallowing), this subclinical inflammation could become chronic — a slow burn that primes the lungs for disease.

With this understanding of the healthy baseline, let us turn to three diseases in which the lung microbiome is dramatically altered.

COPD: A Smouldering Ecosystem

Chronic obstructive pulmonary disease — a progressive, largely irreversible destruction of the airways and alveoli, overwhelmingly caused by cigarette smoking — is the third leading cause of death worldwide. Bacterial infections have long been recognised as central to COPD. Acute exacerbations — episodes of worsening symptoms that punctuate the disease’s chronic course — are frequently triggered by bacterial infection, and the acquisition of new bacterial strains is one of the strongest predictors of an exacerbation. But the role of bacteria in COPD extends beyond acute infections.

Culture-independent studies have revealed that the lung microbiome in COPD is altered in ways that correlate with disease severity. In an early and influential study, Erb-Downward and colleagues at the University of Michigan showed that the lungs of COPD patients had reduced microbial diversity compared to healthy controls, and that this reduction was most pronounced in areas of the lung with the most severe emphysematous destruction. In the most damaged regions, the diverse background community gave way to communities dominated by a single genus — often *Pseudomonas* or *Haemophilus* [REF:han2012].

This pattern — loss of diversity and dominance by a single pathogenic genus — echoes the dysbiosis we encountered in the gut (Chapter 5) and on the skin (Chapter 6). It is the microbial signature of a degraded ecosystem, and in COPD it has clinical consequences. Patients whose lung microbiomes show reduced diversity experience more frequent exacerbations, faster decline in lung function, and greater inflammation. The microbiome composition also shifts in response to treatment: inhaled corticosteroids, the mainstay of COPD management, alter the lung microbial community — though whether these changes are beneficial or harmful remains unclear [REF:mammen2016].

The ecological model proposed by Dickson applies directly to COPD. Smoking damages the mucociliary escalator, reducing the lung's ability to clear aspirated organisms. Chronic inflammation alters the local environment — increasing mucus production, changing the nutrient landscape, and creating oxygen gradients that favour certain organisms over others. The result is a progressive shift from an immigration-dominated community (healthy baseline) to a growth-dominated one (disease state), in which bacteria that would normally be cleared instead establish persistent populations.

Asthma: A Different Kind of Dysbiosis

Asthma presents a different picture. Unlike COPD, asthma is not primarily a disease of structural destruction. It is a disease of airway hyperreactivity — an exaggerated inflammatory response to triggers (allergens, cold air, exercise, infections) that causes the airway smooth muscle to constrict, the airway lining to swell, and mucus production to increase. The result is the episodic breathlessness, wheezing, and chest tightness that characterise the condition.

The first evidence that the lung microbiome might be involved in asthma came from the Hilty study of 2010, which found increased Proteobacteria (especially *Haemophilus*) in the airways of asthmatics compared to controls [REF:hilty2010]. Subsequent studies have confirmed and extended this finding. The asthmatic lung microbiome is characterised by reduced diversity compared to healthy controls, increased relative abundance of Proteobacteria (particularly *Haemophilus influenzae*, *Moraxella catarrhalis*, and *Neisseria* species), and decreased relative abundance of Bacteroidota (particularly *Prevotella*) [REF:chung2017].

What makes the asthma-microbiome relationship particularly interesting is that it appears to operate at multiple levels.

At the **developmental level**, early-life microbial exposures influence asthma risk — a theme we explored in the context of the hygiene hypothesis in Chapter 6. Infants whose nasopharyngeal microbiomes are dominated by *Streptococcus* or *Moraxella* in the first year of life are at increased risk of developing asthma, while those dominated by *Corynebacterium* and *Dolosigranulum* are relatively protected. Early viral infections — particularly RSV — can reshape the nasopharyngeal microbiome in ways that favour Proteobacteria, potentially priming the airway for asthma development [REF:dicicco2018].

At the **clinical level**, the composition of the lung microbiome correlates with asthma phenotype and severity. Patients with more severe, neutrophilic asthma (a subtype characterised by neutrophil-driven rather than eosinophil-driven inflammation) tend to have higher airway bacterial burden and greater enrichment with Proteobacteria. Some studies have found that the microbiome can predict corticosteroid responsiveness: patients with *Haemophilus*-enriched airways may respond less well to inhaled corticosteroids, a finding that has implications for treatment personalisation [REF:chung2017ats].

At the **mechanistic level**, the lung microbiome may directly modulate airway inflammation. The Segal studies described above showed that oral taxa in the lungs drive a Th17-type inflammatory response. In asthma, dysregulated Th17 immunity is implicated in the neutrophilic subtype of the disease. This raises the intriguing possibility that the altered lung microbiome in asthma is not merely a consequence of the inflamed airways but is actively shaping the inflammatory phenotype — a feedback loop in which microbes and inflammation reinforce each other.

Whether manipulation of the airway microbiome could treat or prevent asthma remains an open

question. The idea is tantalising — and it connects directly to the broader story of microbiome therapeutics that we will explore in Part VI. But the evidence base is thin. Clinical trials of probiotics for asthma have been largely disappointing (for reasons that Chapter 23 will explore in detail), and targeted modification of the airway microbiome is technically challenging. We are still in the descriptive phase: mapping the territory, identifying associations, and building the ecological models that will eventually guide intervention.

Cystic Fibrosis: Polymicrobial Complexity

Cystic fibrosis (CF) offers the most dramatic illustration of what happens when the lung's microbial ecosystem collapses. CF is caused by mutations in the gene encoding the **cystic fibrosis transmembrane conductance regulator (CFTR)**, a chloride channel that regulates the viscosity of mucus in the airways, the gut, and other organs. When CFTR is defective, the mucus lining the airways becomes abnormally thick and sticky, clogging the airways, impairing mucociliary clearance, and creating stagnant pools of secretions that are ideal breeding grounds for bacteria.

For decades, the clinical microbiology of CF focused on a small number of traditional pathogens identified by culture: **Staphylococcus aureus** in young children, **Haemophilus influenzae** in early disease, and — most ominously — **Pseudomonas aeruginosa**, whose chronic colonisation of the CF airway is associated with accelerated lung function decline and poor outcomes. The eradication of *Pseudomonas* from the CF airway, or at least the delay of chronic colonisation, became a central goal of CF care.

Culture-independent sequencing revealed that this picture was incomplete. The CF airway is not colonised by one or two pathogens in isolation. It is home to a complex, polymicrobial community that includes obligate anaerobes (*Prevotella*, *Veillonella*, *Fusobacterium*), facultative anaerobes, and organisms rarely detected by routine culture. This community changes over time and across the disease course. In early childhood, the CF airway microbiome is relatively diverse, resembling (though not identical to) the community found in healthy lungs. As the disease progresses and the airways become more damaged, diversity declines progressively, and the community becomes increasingly dominated by one or a few species — usually *Pseudomonas aeruginosa*, but sometimes *Burkholderia cepacia* complex, *Stenotrophomonas*, or *Achromobacter* [REF:frayman2017].

This trajectory — from diverse to monodominant — mirrors the ecological degradation seen in COPD and in the gut during severe illness. It is, once again, the signature of an ecosystem in collapse. And it correlates with clinical outcomes: patients whose airway communities retain diversity tend to have better lung function than those whose communities have collapsed to a single dominant species [REF:han2012].

The discovery of the polymicrobial nature of CF lung infection has practical implications. Anaerobic bacteria in the CF airway, invisible to standard aerobic culture, may contribute to inflammation and tissue damage. They may also interact with the traditional pathogens in ways that affect treatment response — for instance, certain anaerobes produce short-chain fatty acids that can promote *Pseudomonas* biofilm formation, making it harder to eradicate. Understanding these inter-species interactions is essential if we are to move beyond the current “one pathogen, one antibiotic” model of CF treatment toward a more ecological approach.

Where This Matters: The lung function of CF patients has improved dramatically over the past two decades, thanks to aggressive early treatment and, most recently, to CFTR modulator therapies (drugs like ivacaftor and elexacaftor/tezacaftor/ivacaftor)

that partially restore CFTR function. These modulators have been transformative — but their effects on the airway microbiome are only beginning to be understood. Early studies suggest that CFTR modulators shift the airway community toward a more diverse, less pathogen-dominated state. If confirmed, this would provide direct evidence that restoring the physical environment of the lungs (less viscous mucus, better clearance) allows the microbial ecosystem to recover — a powerful example of the link between ecology and therapy.

7.4 The Lung Mycobiome: *Aspergillus* in Cystic Fibrosis and Post-COVID Fungal Superinfections

Just as the gut and the skin harbour fungi alongside bacteria, so does the respiratory tract. And just as elsewhere in the body, the fungal community of the lungs — the **lung mycobiome** — has been far less studied than the bacterial community, for the familiar methodological reasons discussed in Chapter 4: fungi are harder to lyse, harder to sequence, and less well represented in reference databases.

What we know, however, is enough to warrant attention.

In healthy lungs, fungi are present at very low levels, and the species detected are a mixture of environmental organisms (inhaled spores from *Aspergillus*, *Cladosporium*, *Penicillium*, and other common moulds) and organisms carried from the oral cavity (*Candida* species). Their presence is transient and clinically insignificant in immunocompetent individuals. The lung's defences — alveolar macrophages, neutrophils, epithelial cell signalling — are exquisitely effective at neutralising inhaled fungal spores before they can germinate and establish.

But when those defences are compromised, fungi become dangerous. And no genus illustrates this more vividly than ***Aspergillus***.

Aspergillus is ubiquitous. Its spores are everywhere — in soil, in decaying vegetation, in air conditioning systems, in hospital wards. A healthy person inhales hundreds of *Aspergillus* spores every day without consequence. But in patients whose lungs are structurally damaged (as in CF or bronchiectasis), or whose immune systems are suppressed (as in organ transplant recipients, patients on high-dose corticosteroids, or those with haematological malignancies), *Aspergillus* can cause a spectrum of disease that ranges from allergic sensitisation to fatal invasive infection.

In **cystic fibrosis**, *Aspergillus fumigatus* is the most commonly isolated fungal pathogen, detected in the sputum of 10 to 60 per cent of patients depending on the study and the detection method used. The thick, stagnant mucus of the CF airway creates an environment where inhaled spores are not cleared efficiently. Instead, they can germinate, producing hyphae (filamentous growth forms) that penetrate the mucus and interact with both the airway epithelium and the immune system. The clinical consequences range across a spectrum. At one end is **allergic bronchopulmonary aspergillosis (ABPA)** — an exaggerated allergic immune response to *Aspergillus* antigens that causes airway inflammation, mucus plugging, and progressive lung damage. At the other end is ***Aspergillus* bronchitis** — a chronic infection of the airways without tissue invasion, characterised by persistent cough, sputum production, and gradual functional decline [REF:cuthbertson2020].

Culture-independent studies using ITS sequencing (the fungal equivalent of 16S bacterial sequencing) have shown that the CF lung mycobiome is more complex than culture alone suggested. Along-

side *Aspergillus*, species from the genera *Candida*, *Scedosporium*, and *Exophiala* are frequently detected, sometimes as dominant members of the community. CF patients with overt fungal disease show distinct mycobiome profiles dominated by one or a few genera, while those without clinical fungal disease have more diverse fungal communities — echoing the bacterial pattern of diversity loss during disease progression [REF:cuthbertson2020].

The interactions between fungi and bacteria in the CF airway are complex and clinically important. *Aspergillus* and *Pseudomonas aeruginosa* — the two most important pathogens in CF — engage in a competitive relationship. *Pseudomonas* produces small molecules that can inhibit *Aspergillus* germination, and some studies have found an inverse relationship between *Pseudomonas* and *Aspergillus* abundance in the CF airway. Conversely, aggressive antibiotic treatment directed against *Pseudomonas* may inadvertently promote fungal overgrowth by removing bacterial competitors — a scenario directly analogous to the *Candida* blooms seen in the gut after broad-spectrum antibiotics, which we will explore in Chapter 19.

The most dramatic demonstration of fungal opportunism in the respiratory tract, however, came not from CF but from a pandemic.

Post-COVID Fungal Superinfections: A Perfect Storm

When SARS-CoV-2 swept the world beginning in late 2019, it brought with it a wave of secondary infections that caught clinicians off guard. Among the most alarming were fungal superinfections — conditions that had previously been rare outside of specific high-risk populations but that emerged with striking frequency in severe COVID-19 patients.

The first to attract attention was **COVID-associated pulmonary aspergillosis (CAPA)**. Intensive care physicians in Europe and China began reporting invasive *Aspergillus* infections in mechanically ventilated COVID-19 patients at rates far exceeding what was seen in typical ICU populations. Estimates varied, but studies from the Netherlands, Belgium, and France reported CAPA rates of 20 to 35 per cent among ICU patients with COVID-19 — figures that were shocking even to infectious disease specialists accustomed to opportunistic infections in the immunocompromised [REF:jenks2021].

CAPA was devastating. Patients were already critically ill from the virus. Their lungs were inflamed, their alveoli damaged, their immune systems dysregulated by the “cytokine storm” that characterised severe COVID-19. The treatments used to control this inflammation — high-dose corticosteroids (dexamethasone, which became standard of care after the RECOVERY trial in 2020) and IL-6 inhibitors (tocilizumab) — were essential for survival but had the side effect of further suppressing antifungal immunity. The virus damaged the airway epithelium, disrupting the physical barrier that normally prevents *Aspergillus* spores from germinating. And mechanical ventilation bypassed the upper airway defences entirely, providing a direct route for environmental fungi to reach the damaged lungs [REF:worku2022].

CAPA was a convergence of ecological catastrophes: viral destruction of the lung epithelium, immune suppression by both the disease and its treatment, and breach of the physical barriers that normally keep fungal spores in check. It was, in microbial ecological terms, a niche vacancy event followed by an invasion — exactly the kind of scenario that Dickson’s adapted island model would predict.

But CAPA was not the only fungal threat. In India, the second COVID-19 wave in the spring of 2021 triggered an unprecedented outbreak of **mucormycosis** — a rapidly progressive, often fatal

infection caused by fungi of the order Mucorales (particularly *Rhizopus*, *Mucor*, and *Lichtheimia*). India reported over 45,000 cases of COVID-associated mucormycosis (CAM) in a matter of months, primarily manifesting as rhino-orbito-cerebral mucormycosis — an infection that begins in the sinuses and can spread to the orbits and brain with terrifying speed.

The Indian outbreak was driven by a specific combination of risk factors that converged in that population: a high prevalence of poorly controlled diabetes mellitus (uncontrolled hyperglycaemia provides Mucorales with the iron and glucose they need to grow), widespread and often indiscriminate use of corticosteroids (which suppress immune defences and raise blood glucose further), and a baseline mucormycosis incidence that was already higher in India than in most other countries. The COVID-19 pandemic did not create the underlying susceptibility — it amplified it. Viral lung damage, corticosteroid-induced immune suppression, and diabetic metabolic dysfunction combined to create conditions in which Mucorales fungi, normally held in check by intact immunity, could invade tissue with devastating consequences [REF:pal2021].

The mortality rate was appalling: approximately 34 per cent of reported cases, with higher rates in those with cerebral involvement. Survivors frequently required extensive surgery, including orbital exenteration (removal of the eye and surrounding tissue) — a price paid for delayed diagnosis of a condition that many clinicians had never seen.

The post-COVID fungal superinfections are a case study in microbial ecology under extreme perturbation. They demonstrated, with tragic clarity, that the respiratory ecosystem's defences are layered and interdependent: remove the physical barrier (viral epithelial damage), suppress the immune surveillance (corticosteroids, cytokine storm), and alter the metabolic environment (hyperglycaemia), and organisms that are normally inconsequential become lethal. The fungi did not change. The ecosystem they inhabited did.

7.5 Lessons from COVID-19: Viral Infection and Microbial Disruption

The COVID-19 pandemic was a global catastrophe. It was also, unwittingly, the largest natural experiment ever conducted on the interaction between a viral infection and the respiratory microbiome.

SARS-CoV-2 did not simply infect and kill lung cells. It restructured the entire microbial ecology of the respiratory tract — and of the gut, too, through mechanisms we are still working to understand. The virus's effects on the respiratory microbiome were multifaceted, operating through direct tissue damage, immune dysregulation, and the consequences of the treatments used to manage it.

The first and most direct effect was **epithelial destruction**. SARS-CoV-2 enters cells via the ACE2 receptor, which is expressed on the surface of the airway epithelial cells that line the nose, throat, and lungs. The virus replicates within these cells and kills them, disrupting the physical barrier that separates the microbial world from the body's internal tissues. This is not unique to COVID-19 — influenza and other respiratory viruses cause similar damage — but the extent of epithelial injury in severe COVID-19 was exceptional, particularly in the alveoli, where it manifested as diffuse alveolar damage, the pathological hallmark of acute respiratory distress syndrome (ARDS).

The destruction of the epithelial barrier has cascading ecological consequences. The mucociliary escalator is disrupted, reducing microbial clearance. The denuded epithelial surface provides new attachment sites for bacteria. Serum proteins and cellular debris leak into the airway lumen, provid-

ing nutrients that promote bacterial growth. And the immune response — initially suppressed in the early phase of severe COVID-19, then explosively overactivated in the cytokine storm phase — creates an inflammatory environment that selects for organisms adapted to inflamed tissue.

The second effect was on the **immune response** itself. Severe COVID-19 is characterised by a paradoxical immune dysfunction: impaired adaptive immunity (reduced T cell counts, impaired interferon responses) combined with excessive innate immune activation (the cytokine storm, characterised by high levels of IL-6, TNF- α , and other pro-inflammatory mediators). This combination is ecologically catastrophic. The impaired adaptive response weakens the targeted killing of specific pathogens, while the excessive innate response creates collateral tissue damage that opens new niches for microbial exploitation.

A remarkable study by Tsitsiklis, Langelier, and colleagues at the University of California San Francisco illustrated this dynamic in detail. They followed 28 COVID-19 patients in the ICU, 15 of whom developed ventilator-associated pneumonia (VAP). Using a combination of bulk RNA sequencing and single-cell RNA sequencing of tracheal aspirate samples, they found that COVID-19 patients who subsequently developed VAP showed a striking impairment in immune signalling — detectable a full two weeks before the secondary infection became clinically apparent. Longitudinal metatranscriptomic analysis (profiling which microbial genes were actively being expressed, not just which organisms were present) revealed progressive disruption of the lung microbiome community composition in the pre-VAP period, with outgrowth of opportunistic pathogens that correlated with the impaired immune signalling [REF:tsitsiklis2021].

This finding is remarkable for several reasons. It shows that secondary bacterial pneumonia in COVID-19 is not a sudden event but the endpoint of a process that unfolds over days to weeks. It demonstrates that the immune impairment and the microbial disruption are linked — that the failing immune defence allows the microbiome to shift toward pathogenic dominance. And it suggests that there may be a window of opportunity — a pre-clinical phase during which the trajectory toward secondary infection is detectable and, potentially, modifiable.

The third effect was **treatment-mediated**. Dexamethasone, the corticosteroid shown by the RECOVERY trial to reduce mortality in severe COVID-19, is a powerful immunosuppressant. It was lifesaving for many patients, but it also suppressed the antifungal and antibacterial immune responses, contributing to the CAPA and mucormycosis outbreaks described in the previous section. Broad-spectrum antibiotics, administered prophylactically or empirically to COVID-19 patients in many settings, disrupted the bacterial microbiome and may have promoted fungal overgrowth. And mechanical ventilation — necessary for patients who could not breathe on their own — bypassed the natural airway defences, creating a direct conduit for environmental and oropharyngeal organisms to reach the damaged lungs.

The COVID-19 experience reinforced a lesson that microbiome science has been teaching for two decades: the respiratory ecosystem is not a collection of isolated infections waiting to happen. It is an integrated system in which the physical barriers, the immune response, the microbial communities, and the treatments we administer all interact. Disrupt one component, and the ripple effects propagate through the entire system.

Beyond the lungs, COVID-19 also affected the **gut microbiome**. Several studies, most notably by Yeoh and colleagues in Hong Kong, showed that patients with COVID-19 had altered gut microbial communities compared to controls, with reduced abundance of beneficial commensals (including *Faecalibacterium prausnitzii* and *Bifidobacterium* species) and increased abundance of opportunistic pathogens. The degree of gut dysbiosis correlated with disease severity and with inflammatory

markers in the blood, raising the question — still unresolved — of whether the gut-lung axis (the bidirectional communication between the gut microbiome and the respiratory immune system, which we will explore in greater depth in Chapter 12) played a role in determining COVID-19 outcomes.

The pandemic also exposed the fragility of ecological thinking in clinical practice. In the rush to save critically ill patients, antibiotics were prescribed widely — often before bacterial infection was confirmed, simply because clinicians could not afford to wait for culture results. This was understandable and often necessary. But it meant that many patients received broad-spectrum antibiotics that disrupted their microbiomes without clear evidence that bacterial infection was present. The tension between the urgency of acute care and the ecological consequences of treatment is not new — it is the central dilemma of antibiotic stewardship, which we will address in Chapter 19 — but COVID-19 brought it into sharper relief than any previous crisis.

Where This Matters: If you have ever been admitted to hospital with a severe respiratory infection, or if you care for patients who are, the lessons of COVID-19's microbial ecology are directly relevant. Secondary infections — bacterial and fungal — are not random misfortunes. They are ecological consequences of the disruptions caused by the primary infection and its treatment. Recognising this means looking beyond the immediate pathogen and thinking about the ecosystem: the barriers, the immune responses, the microbial communities, and the drugs that affect all three. The best clinicians already think this way. The microbiome gives them a language for it.

In the next chapter, we move to a body site where the relationship between microbes and health is more intimate still — the urogenital tract, where the microbiome shapes everything from fertility to preterm birth, and where one of the most distinctive microbial communities in the human body quietly maintains a partnership that has lasted for millennia.

Chapter 7 — Key References

- [REF:hilty2010] Hilty, M., Burke, C., Pedro, H., Cardenas, P., Bush, A., Bossley, C., ... & Cookson, W.O.C. (2010). Disordered microbial communities in asthmatic airways. *PLoS One*, 5(1), e8578. DOI: 10.1371/journal.pone.0008578
- [REF:dickson2015a] Dickson, R.P., Erb-Downward, J.R. & Huffnagle, G.B. (2015). Homeostasis and its disruption in the lung microbiome. *American Journal of Physiology — Lung Cellular and Molecular Physiology*, 309(10), L1047–L1055. DOI: 10.1152/ajplung.00279.2015
- [REF:dickson2015b] Dickson, R.P. (2016). The microbiome and critical illness. *Lancet Respiratory Medicine*, 4(1), 59–72. DOI: 10.1016/S2213-2600(15)00427-0
- [REF:dickson2016] Dickson, R.P., Erb-Downward, J.R., Martinez, F.J. & Huffnagle, G.B. (2016). The microbiome and the respiratory tract. *Annual Review of Physiology*, 78, 481–504. DOI: 10.1146/annurev-physiol-021115-105238
- [REF:segal2016] Segal, L.N., Clemente, J.C., Tsay, J.C.J., Koralov, S.B., Keller, B.C., Wu, B.G., ... & Weiden, M.D. (2016). Enrichment of the lung microbiome with oral taxa is associated with lung inflammation of a Th17 phenotype. *Nature Microbiology*, 1, 16031. DOI: 10.1038/nmicrobiol.2016.31

- [REF:cleary2017] Cleary, D.W. & Clarke, S.C. (2017). The nasopharyngeal microbiome. *Emerging Topics in Life Sciences*, 1(4), 297–312. DOI: 10.1042/ETLS20170041
- [REF:desteenhuijsen2016] de Steenhuijsen Piter, W.A.A., Heinonen, S., Hasrat, R., Bun-sow, E., Smith, B., Suarez-Arrabal, M.C., ... & Mejias, A. (2016). Nasopharyngeal microbiota, host transcriptome, and disease severity in children with respiratory syncytial virus infection. *American Journal of Respiratory and Critical Care Medicine*, 194(9), 1104–1115. DOI: 10.1164/rccm.201602-0220OC
- [REF:mammen2016] Mammen, M.J. & Sethi, S. (2016). COPD and the microbiome. *Respirology*, 21(4), 590–599. DOI: 10.1111/resp.12732
- [REF:han2012] Han, M.K., Huang, Y.J., Lipuma, J.J., Boushey, H.A., Boucher, R.C., Cookson, W.O., ... & Martinez, F.J. (2012). Significance of the microbiome in obstructive lung disease. *Thorax*, 67(5), 456–463. DOI: 10.1136/thoraxjnl-2011-201183
- [REF:chung2017] Chung, K.F. (2017). Airway microbial dysbiosis in asthmatic patients: A target for prevention and treatment? *Journal of Allergy and Clinical Immunology*, 139(4), 1071–1081. DOI: 10.1016/j.jaci.2017.02.004
- [REF:chung2017ats] Chung, K.F. (2017). Potential role of the lung microbiome in shaping asthma phenotypes. *Annals of the American Thoracic Society*, 14(Suppl 5), S326–S331. DOI: 10.1513/AnnalsATS.201702-138AW
- [REF:dicicco2018] Di Cicco, M., Pistello, M., Jacinto, T., Ragazzo, V., Piras, M., Freer, G., Pifferi, M. & Peroni, D. (2018). Does lung microbiome play a causal or casual role in asthma? *Pediatric Pulmonology*, 53(10), 1340–1345. DOI: 10.1002/ppul.24086
- [REF:frayman2017] Frayman, K.B., Armstrong, D.S., Grimwood, K. & Ranganathan, S.C. (2017). The airway microbiota in early cystic fibrosis lung disease. *Pediatric Pulmonology*, 52(11), 1384–1404. DOI: 10.1002/ppul.23782
- [REF:cuthbertson2020] Cuthbertson, L., Felton, I., James, P., Cox, M.J., Bilton, D., Schelenz, S., ... & Moffatt, M.F. (2020). The fungal airway microbiome in cystic fibrosis and non-cystic fibrosis bronchiectasis. *Journal of Cystic Fibrosis*, 20(2), 295–302. DOI: 10.1016/j.jcf.2020.05.013
- [REF:jenks2021] Jenks, J.D., Nam, H.H. & Hoenigl, M. (2021). Invasive aspergillosis in critically ill patients: Review of definitions and diagnostic approaches. *Mycoses*, 64(9), 1002–1014. DOI: 10.1111/myc.13274
- [REF:worku2022] Worku, D.A. (2022). SARS-CoV-2 associated immune dysregulation and COVID-associated pulmonary aspergillosis (CAPA): A cautionary tale. *International Journal of Molecular Sciences*, 23(6), 3228. DOI: 10.3390/ijms23063228
- [REF:pal2021] Pal, R., Singh, B., Bhadada, S.K., Banerjee, M., Bhogal, R.S., Hage, N. & Kumar, A. (2021). COVID-19-associated mucormycosis: An updated systematic review of literature. *Mycoses*, 64(12), 1452–1459. DOI: 10.1111/myc.13338
- [REF:tsitsiklis2021] Tsitsiklis, A., Zha, B.S., Byrne, A., Devoe, C., Levan, S., Rackaityte, E., ... & Langelier, C.R. (2021). Impaired immune signaling and changes in the lung microbiome precede secondary bacterial pneumonia in COVID-19. *Research Square*. DOI: 10.21203/rs.3.rs-380803/v1

Additional references to be verified and expanded during manuscript preparation.

Chapter 8: The Urogenital Microbiome

In the previous chapter, we followed the respiratory tract from the nasopharynx to the deepest alveoli, watching the “sterile lung” dogma crumble under the weight of culture-independent evidence. Now we turn to another body system where the assumption of sterility persisted for even longer — and where its dismantling has had equally profound consequences for how we understand health and disease.

The urogenital tract occupies a unique position in the human body. It is simultaneously an excretory system, a reproductive system, and a microbial habitat of remarkable specialisation. For women, the vagina harbours one of the most distinctive microbial communities found anywhere in the human body — an ecosystem dominated by a single genus, *Lactobacillus*, to a degree unmatched at any other body site. For both sexes, the urinary tract was long considered a sterile conduit — a plumbing system through which urine flowed from kidney to bladder to outside world without encountering a single living microbe. That assumption, like the sterile lung dogma, was wrong.

This chapter explores the microbial ecology of the urogenital tract in four parts. We begin with the vaginal microbiome — its unusual *Lactobacillus* dominance, the community state type framework that has transformed how we classify it, and the surprising lesson that “normal” does not look the same in all women. We then turn to the urinary microbiome — the **urobiome** — and the story of how improved culture techniques and DNA sequencing revealed a microbial community in what was supposed to be sterile urine. From there, we examine the implications of urogenital microbial communities for fertility, pregnancy, and preterm birth — one of the most clinically urgent areas in microbiome research. Finally, we consider the broader question of sex differences in microbiome composition and what they mean for disease risk.

8.1 The Vaginal Microbiome — *Lactobacillus* Dominance and Community State Types

The vagina is, microbiologically speaking, an outlier.

At most body sites, a healthy microbiome is characterised by diversity. The gut harbours hundreds of species across multiple phyla. The skin hosts distinct communities of bacteria, fungi, and viruses that vary from one anatomical niche to the next. Even the sparse lung microbiome reflects a broad spectrum of organisms washed down from the oral cavity. Ecological diversity, as we have seen throughout this book, is generally associated with stability and resilience — a diverse community is harder to disrupt and quicker to recover.

The vagina defies this pattern. In the majority of reproductive-age women, the vaginal microbiome is dominated — often overwhelmingly — by a single genus: **Lactobacillus**. It is not unusual for *Lactobacillus* species to account for 90 per cent or more of the total bacterial community. This degree of dominance by one genus is remarkable. It is also, as we shall see, highly functional.

The story of how we came to understand vaginal microbial communities in their current detail begins with a landmark study published in 2011 by Jacques Ravel, Larry Forney, and colleagues at the

University of Maryland. Using pyrosequencing of 16S rRNA genes — the same molecular barcode approach that had revealed the complexity of the gut and skin microbiomes — they characterised the vaginal bacterial communities of 396 asymptomatic, reproductive-age North American women drawn from four ethnic groups: white, Black, Hispanic, and Asian [REF:ravel2011].

What Ravel and colleagues found was both elegant and surprising. The vaginal communities did not exist on a smooth continuum. Instead, they clustered into five distinct groups, which the researchers designated **community state types** (CSTs). Four of the five CSTs were each dominated by a different *Lactobacillus* species:

- **CST I:** dominated by *Lactobacillus crispatus*
- **CST II:** dominated by *Lactobacillus gasseri*
- **CST III:** dominated by *Lactobacillus iners*
- **CST V:** dominated by *Lactobacillus jensenii*

The fifth group — **CST IV** — was different. It lacked *Lactobacillus* dominance altogether and instead harboured a diverse, polymicrobial community of strict and facultative anaerobes, including *Gardnerella*, *Prevotella*, *Atopobium*, *Megasphaera*, and *Sneathia*. This community profile resembled what clinicians had long recognised as the microbial signature of **bacterial vaginosis** (BV) — and yet many of the women with CST IV communities were asymptomatic. They had no discharge, no odour, no complaints. By every clinical measure they considered themselves healthy [REF:ravel2011].

This finding was important for several reasons. It established a classification framework — the CST system — that gave researchers a common language for describing vaginal microbial communities. It demonstrated that the vaginal microbiome, like the gut enterotypes described around the same time (Chapter 5), falls into recognisable configurations rather than varying randomly from person to person. And it raised a provocative question: if some healthy women naturally harbour non-*Lactobacillus* communities, who defines what “normal” is?

The Four *Lactobacilli* — Not All Equal

The CST framework also revealed that not all *Lactobacillus* species are equivalent. Although all four *Lactobacillus*-dominated CSTs share the key ecological function of producing **lactic acid** — which maintains the vaginal pH between 3.5 and 4.5, creating an environment inhospitable to most pathogens — the species differ in important ways.

Lactobacillus crispatus (CST I) is widely considered the most protective species. It produces both D- and L-isomers of lactic acid, generates hydrogen peroxide under aerobic conditions, and maintains the lowest and most stable vaginal pH. Women whose vaginal microbiomes are dominated by *L. crispatus* have the lowest rates of bacterial vaginosis, sexually transmitted infections (including HIV, chlamydia, and gonorrhoea), and adverse pregnancy outcomes. CST I communities are also the most temporally stable — once established, they tend to persist, resisting perturbation from menstruation, sexual activity, and minor antibiotic exposures [REF:smith2017].

Lactobacillus iners (CST III) presents a more ambiguous picture. It is the most commonly detected *Lactobacillus* species across all ethnic groups and is found both in healthy women and in women with or transitioning to bacterial vaginosis. Unlike *L. crispatus*, *L. iners* produces only the L-isomer of lactic acid (which is less potently antimicrobial than the D-isomer), does not produce hydrogen peroxide, and has a smaller genome with fewer genes encoding antimicrobial compounds. Its communities tend to be less stable over time, more susceptible to transitioning to the diverse

CST IV profile, and — as we shall see in section 8.3 — associated with a higher risk of preterm birth. Some researchers have described *L. iners* as a “transitional” species — one that can coexist with either a *Lactobacillus*-dominated or a diverse anaerobic community, but that does not confer the same degree of protection as *L. crispatus* [REF:smith2017].

Lactobacillus gasseri (CST II) and *Lactobacillus jensenii* (CST V) are less common and less extensively studied, but they appear to fall somewhere between *L. crispatus* and *L. iners* in terms of their protective effects.

Why Lactic Acid Matters

The ecological strategy of vaginal *Lactobacillus* species is elegantly simple: they ferment **glycogen**. The vaginal epithelium of reproductive-age women, under the influence of oestrogen, is rich in glycogen — a complex sugar that is deposited in the epithelial cells and released as they are shed. *Lactobacillus* species metabolise this glycogen (and its breakdown products, particularly maltose) to produce lactic acid, which drives the vaginal pH down to 3.5–4.5.

This acidic environment is directly inhibitory to most pathogenic bacteria. But the protection extends beyond pH alone. Lactic acid, particularly the D-isomer produced abundantly by *L. crispatus*, has direct antimicrobial properties. It disrupts the outer membranes of gram-negative bacteria, inhibits the growth of *Chlamydia trachomatis*, *Neisseria gonorrhoeae*, and HIV in laboratory studies, and potentiates the antimicrobial activity of other innate immune molecules. In effect, *Lactobacillus* species create a **chemical shield** — an acidic, antimicrobial environment maintained by metabolism rather than by the immune system.

This is a partnership of remarkable specificity. The host provides the substrate (glycogen, via oestrogen-driven epithelial maturation). The bacteria convert it to an antimicrobial product (lactic acid). The result is an environment hostile to pathogens but comfortable for the *Lactobacillus* species themselves, which have evolved to thrive at low pH. It is mutualism in its purest form — and it explains why the vaginal microbiome is so sensitive to changes in oestrogen levels, a point we will return to shortly.

Ethnic Variation and the Problem of “Normal”

One of the most striking findings of the Ravel study was that the distribution of CSTs varied significantly across ethnic groups. Among white and Asian women, approximately 80–90 per cent had *Lactobacillus*-dominated communities (CSTs I–III and V). Among Hispanic and Black women, the proportion was substantially lower — roughly 60–70 per cent — with a correspondingly higher prevalence of the diverse, non-*Lactobacillus* CST IV profile [REF:ravel2011].

This finding had immediate clinical implications. Bacterial vaginosis is traditionally diagnosed using criteria that assume a *Lactobacillus*-dominant community is normal and a diverse anaerobic community is pathological. The Nugent scoring system, the laboratory gold standard for BV diagnosis, assigns a score based on the relative proportions of *Lactobacillus*-like, *Gardnerella*-like, and curved gram-negative rod morphotypes on a Gram-stained vaginal smear. A high Nugent score — indicating few *Lactobacillus* and abundant diverse anaerobes — equates to a BV diagnosis.

But if a substantial proportion of healthy women — particularly Black and Hispanic women — naturally harbour CST IV communities without any symptoms, then the Nugent score is diagnosing a microbiological state, not necessarily a disease. The clinical significance of this observation remains debated, but it has prompted a re-evaluation of what constitutes a “healthy” vaginal microbiome

and whether diagnostic thresholds calibrated primarily on white populations can be applied universally.

The reasons for ethnic variation in vaginal CST distribution are not fully understood. Genetic differences in the immune response, variation in glycogen metabolism, differences in the vaginal pH set point, and socioeconomic and behavioural factors (diet, sexual practices, hygiene products) have all been proposed. What is clear is that the relationship between vaginal microbial community composition and health is more nuanced than the simple binary of “*Lactobacillus* = good, everything else = bad.”

Temporal Dynamics — A Community in Motion

The vaginal microbiome is not static. Studies following individual women over time — most notably the longitudinal study by Gajer, Ravel, and colleagues in 2012 — have shown that vaginal community composition fluctuates over days and weeks, with some women maintaining stable CST profiles for months and others transitioning between CSTs multiple times per menstrual cycle [REF:gajer2012].

Menstruation is a particularly significant perturbation. The influx of blood — which has a neutral pH and contains iron, proteins, and other nutrients — temporarily disrupts the acidic environment maintained by *Lactobacillus*. In some women, this leads to a transient shift from a *Lactobacillus*-dominated community to a more diverse profile, which then reverts as menstruation ends and the acidic environment is re-established. In others, the community remains stable throughout the cycle.

Sexual intercourse introduces alkaline seminal fluid (pH 7.0–8.0) and exogenous bacteria, both of which can transiently alter vaginal pH and community composition. Antibiotic use — particularly broad-spectrum antibiotics for non-gynaecological indications — can deplete *Lactobacillus* populations and permit overgrowth of anaerobic organisms, sometimes triggering clinical BV.

The degree of temporal stability varies between CSTs. As noted above, CST I (*L. crispatus*-dominated) communities tend to be the most resilient, returning rapidly to their baseline after perturbation. CST III (*L. iners*-dominated) communities are more labile, more frequently transitioning to CST IV and back. This temporal instability may be part of the reason why *L. iners* dominance is associated with higher rates of BV and adverse reproductive outcomes — not because *L. iners* is directly harmful, but because the community it anchors is more easily disrupted.

The Hormonal Connection

The vaginal microbiome is inextricably tied to the hormonal status of the host. Before puberty, when oestrogen levels are low, the vaginal epithelium is thin and glycogen-poor. The vaginal pH is neutral or slightly alkaline, and the microbial community is diverse, with low abundances of *Lactobacillus*. At puberty, rising oestrogen levels stimulate epithelial thickening and glycogen deposition. *Lactobacillus* species, newly supplied with their preferred substrate, expand rapidly and acidify the environment. The adult vaginal microbiome is established.

This partnership continues through the reproductive years, fluctuating with the menstrual cycle (oestrogen peaks at ovulation), shifting dramatically during pregnancy (when sustained high oestrogen levels promote *Lactobacillus* dominance and community stability — a point we will explore in section 8.3), and then unwinding at menopause.

After menopause, as ovarian oestrogen production declines, the vaginal epithelium thins, glycogen levels fall, and the acidic, *Lactobacillus*-dominated community gives way to a more diverse, higher-pH microbial profile reminiscent of the prepubertal state. This transition is associated with an increased risk of urinary tract infections, vaginal dryness, and vulvovaginal atrophy — conditions collectively known as the **genitourinary syndrome of menopause (GSM)**. Vaginal oestrogen therapy can partially reverse these changes, restoring epithelial thickness, glycogen deposition, and *Lactobacillus* dominance — a therapeutic intervention that acts, in essence, by restoring the substrate that feeds the mutualistic microbial community [REF:thomas-white2020].

Where This Matters: If you have ever been prescribed vaginal oestrogen cream after menopause, the rationale is not merely to thicken the vaginal lining. It is to re-establish the conditions under which *Lactobacillus* can thrive — restoring the acidic, antimicrobial environment that these bacteria have maintained since puberty. The cream feeds the epithelium; the epithelium feeds the bacteria; the bacteria defend the ecosystem. You are not treating a hormone deficiency. You are restoring an ecology.

8.2 The Urinary Microbiome — Another “Sterile” Myth Debunked

For well over a century, a foundational principle of clinical microbiology held that urine in the healthy bladder is sterile. The logic seemed straightforward: urine is produced by the kidneys through ultrafiltration of the blood, flows through the ureters into the bladder, and is expelled through the urethra. No bacteria should be present. When bacteria are detected in a urine culture, the reasoning went, they signify either contamination during collection or a urinary tract infection.

This principle was codified in the **standard urine culture (SUC)** — a diagnostic test that has remained virtually unchanged since the 1950s. The standard protocol involves inoculating a urine sample onto a limited number of agar media, incubating it aerobically at 35°C for 24 hours, and reporting anything that grows above a threshold of 10⁵ colony-forming units per millilitre (CFU/mL). If nothing grows — or nothing grows above the threshold — the result is reported as “no growth” or “negative.” Patient and clinician alike interpret this as confirmation: the urine is sterile.

But the standard urine culture was never designed to detect all bacteria. It was designed to detect the most common uropathogens — principally *Escherichia coli* — under growth conditions that favour them. The restricted media, the aerobic incubation, and the short incubation period mean that slow-growing organisms, anaerobes, organisms with specific nutritional requirements, and organisms present at lower concentrations are systematically missed. The test does not prove that urine is sterile. It proves that no organism grew under the specific conditions tested. These are not the same thing.

The challenge to urinary sterility came from two directions simultaneously: improved culture methods and culture-independent DNA sequencing.

The Expanded Culture Revolution

The key breakthrough on the culture side came from the laboratory of Alan Wolfe and Linda Brubaker at Loyola University Chicago. Beginning around 2012, their team developed the **Expanded Quantitative Urine Culture (EQUC)** protocol — a systematic enhancement of the standard urine culture that used a larger volume of urine inoculum, a broader array of culture media (including anaerobic and CO₂-enriched conditions), extended incubation times (up to 48

hours), and lower detection thresholds. The critical difference: EQUUC was designed to find what the standard culture missed.

The results were unequivocal. When catheterised urine samples from women with no symptoms of urinary tract infection were processed by both standard culture and EQUUC, the standard culture reported “no growth” in the majority of cases. EQUUC, applied to the same samples, detected bacteria in the majority of those “sterile” specimens. The organisms found were not exotic or unexpected — they included *Lactobacillus*, *Gardnerella*, *Corynebacterium*, *Streptococcus*, *Staphylococcus*, and *Actinomyces*, among others — but they were present at concentrations and under growth conditions that the standard culture was blind to [REF:price2020].

The parallel application of 16S rRNA gene sequencing to the same catheterised urine samples confirmed and extended the EQUUC findings. DNA-based methods detected bacterial sequences in nearly all samples, including those reported as sterile by both standard culture and EQUUC. The urinary microbiome — soon dubbed the **urobiome** — was real.

What Lives in the Bladder?

The characterisation of the urobiome in healthy women has now been replicated by multiple independent groups. The picture that emerges is of a low-biomass community — far sparser than the gut or the vagina — but a genuine community nonetheless.

The most commonly detected genus in the healthy female bladder is **Lactobacillus**, mirroring its dominance in the vagina and suggesting cross-contamination or shared colonisation between adjacent body sites. Other commonly detected genera include *Gardnerella*, *Streptococcus*, *Staphylococcus*, *Corynebacterium*, and — at lower levels — *Escherichia*. The relative proportions of these genera vary between individuals, and, echoing the vaginal CST framework, researchers have proposed the concept of **urotypes** — distinct community profiles defined by the dominant taxon in each sample. The most common urotype in healthy women is *Lactobacillus*-dominated, but *Gardnerella*-dominated and *Escherichia*-dominated urotypes are also observed [REF:price2020].

Intriguingly, the urobiome composition appears to change with age. In a cross-sectional study of 224 continent women, Price, Hilt, Wolfe, Brubaker, and colleagues found that the *Gardnerella* urotype was more common in younger women, while the *Escherichia* urotype was more common in older women — a shift that likely reflects the declining oestrogen levels and changing vaginal ecology of menopause, which alter the microbial communities of the entire pelvic floor [REF:price2020].

Clinical Implications — Rethinking the “Sterile” Paradigm

The discovery of the urobiome has practical consequences that extend beyond academic interest.

The most immediate implication concerns the standard urine culture itself. A test that reports “no growth” when bacteria are in fact present is not just imprecise — it is potentially misleading. Consider a woman who presents with urinary urgency, frequency, and discomfort. Her urine is cultured using the standard protocol. The result comes back negative. She is told she does not have a urinary tract infection. Her symptoms persist. She returns. The culture is repeated. Negative again. She is told the problem is “functional” or “overactive bladder” and offered anticholinergic medication.

But what if the standard culture simply failed to detect the organism responsible for her symptoms? Wolfe and Brubaker’s work has shown that many women diagnosed with “culture-negative” urinary

symptoms harbour bacteria detectable by EQUIC — organisms like *Aerococcus urinae*, *Actinomyces neuui*, and atypical *Streptococcus* species that are viable, potentially pathogenic, and invisible to the standard test [REF:wolfe2022].

This does not mean that every woman with urinary symptoms and a negative standard culture has an occult infection. The urobiome exists in health as well as disease, and distinguishing a pathological state from normal variation is precisely the challenge that the urobiome discovery has created. But it does mean that the standard culture's binary verdict — “infection” or “sterile” — is a clinical fiction. The bladder, like the lungs before it, has turned out to be not sterile but simply understudied.

The urobiome also raises questions about the long-standing concept of **asymptomatic bacteriuria** — the presence of bacteria in urine without symptoms. Asymptomatic bacteriuria (typically defined as $\geq 10^5$ CFU/mL of a single organism on standard culture) has been a focus of screening and treatment for decades, particularly in pregnant women — where it genuinely increases the risk of pyelonephritis and preterm birth and where treatment is recommended by current guidelines. But in most other populations, including elderly women and catheterised patients, the evidence for treating asymptomatic bacteriuria has progressively weakened, and current guidelines recommend against it. The urobiome adds a new layer to this reassessment: if the bladder normally contains bacteria, the distinction between “asymptomatic bacteriuria” and “a normal urobiome” becomes less clear. Screening may be detecting a community, not an infection — and treating a community with antibiotics is precisely the kind of ecological disruption whose consequences we will explore in Chapter 19.

Where This Matters: The next time you provide a urine sample for culture, it is worth knowing that the test used in most clinical laboratories was designed in the 1950s and has scarcely changed since. It detects a narrow range of organisms under a narrow set of conditions. A “negative” result does not mean your urine is sterile — it means that nothing grew under the conditions tested. This distinction matters most for the millions of women with recurrent urinary symptoms whose cultures are repeatedly negative, and who may be living with treatable infections that the standard test cannot see.

8.3 Implications for Fertility, Pregnancy, and Preterm Birth

If there is one area where the clinical stakes of the urogenital microbiome are highest, it is reproduction. The vaginal microbiome is implicated at every stage — from conception to implantation to the maintenance of pregnancy to the timing of delivery. And of all the adverse outcomes linked to vaginal microbial communities, none has attracted more research attention than **preterm birth** — delivery before 37 weeks of gestation.

Preterm birth is the leading cause of neonatal mortality worldwide and a major cause of long-term disability. Approximately 15 million babies are born preterm each year, and complications of prematurity are the single largest contributor to deaths in children under five. The causes of spontaneous preterm birth are heterogeneous, but infection and inflammation of the reproductive tract have long been recognised as major risk factors. What microbiome research has added to this picture is specificity: not just whether infection is present, but which microbial community the mother carries and how that community interacts with the host immune response.

The Vaginal Microbiome in Pregnancy

Pregnancy reshapes the vaginal microbiome. Under the sustained high oestrogen and progesterone levels of pregnancy, the vaginal epithelium thickens, glycogen production increases, and the microbial community becomes more *Lactobacillus*-dominant, less diverse, and more temporally stable than in the non-pregnant state. This shift — sometimes described as a “locking in” of the community — is thought to be protective, creating an acidic, antimicrobial environment that guards the cervix and uterus against ascending infection during the vulnerable months of gestation [REF:smith2017].

The Integrative Human Microbiome Project (iHMP), the second phase of the NIH Human Microbiome Project, included a major pregnancy and preterm birth component. This landmark effort followed hundreds of pregnant women longitudinally, sampling multiple body sites at multiple time points and integrating microbiome data with host transcriptomics, proteomics, and metabolomics [REF:ihmp2019]. The iHMP and related studies have established several key findings.

First, the vaginal microbiome of women who go on to deliver at term is typically dominated by *Lactobacillus crispatus* (CST I) or, less optimally, by *Lactobacillus iners* (CST III). Women with diverse, *Lactobacillus*-depleted communities (CST IV) during pregnancy are at increased risk of preterm birth.

Second, the timing matters. A systematic review and network meta-analysis by Gudnadottir, Brunselaers, and colleagues in 2022, synthesising data from 17 longitudinal studies, found that women with low-*Lactobacillus* vaginal microbiomes had a significantly increased odds of preterm birth (OR 1.69, 95% CI 1.15–2.49) compared to women with *L. crispatus*-dominated communities. The association was strongest for earlier preterm births (before 34 weeks), which are also the most clinically severe [REF:gudnadottir2022].

Third — and this is where the subtlety lies — the distinction between *L. crispatus* and *L. iners* matters for pregnancy outcomes too. A large meta-analysis by Huang, Callahan, Fettweis, and colleagues in 2023, integrating data from 12 cohorts and over 2,000 pregnant women, found that *L. crispatus* was consistently and significantly associated with reduced preterm birth risk across nearly all studies. By contrast, *L. iners* was associated with *increased* risk of preterm birth — placing it on the same side as *Prevotella*, *Gardnerella*, and other diverse anaerobes. The authors found that earlier preterm births (before 32 weeks) were more predictable from vaginal microbiome composition than later preterm births, suggesting that the microbiome’s influence is greatest in the most severe cases [REF:huang2023].

The Mechanism — From Dysbiosis to Inflammation to Preterm Labour

How does a non-*Lactobacillus* vaginal community contribute to preterm birth? The most widely accepted model involves ascending infection and inflammation.

In a *Lactobacillus*-dominated vaginal environment, the cervical mucus plug and the acidic vaginal milieu provide a dual barrier against microbial ascent into the uterus. When this barrier is weakened — by *Lactobacillus* depletion, elevated vaginal pH, or the presence of sialidase-producing anaerobes like *Gardnerella* and *Prevotella* that can degrade cervical mucus — bacteria can ascend through the cervix into the uterine cavity.

Once in the uterus, these organisms trigger an innate immune response. Pattern recognition receptors on the decidual and fetal membrane cells detect bacterial components (lipopolysaccharides, peptidoglycan, flagellin) and activate inflammatory cascades — principally through NF-κB signalling

and the production of pro-inflammatory cytokines including interleukin-1 β (IL-1 β), interleukin-6 (IL-6), and tumour necrosis factor- α (TNF- α). These cytokines, in turn, stimulate the production of **prostaglandins** and **matrix metalloproteinases** (MMPs) that ripen the cervix, weaken the fetal membranes, and initiate uterine contractions — the physiological cascade of labour, triggered prematurely by infection.

This model — **ascending infection** □ **intrauterine inflammation** □ **premature activation of labour** — is well supported by evidence from both animal models and human studies. It explains why bacterial vaginosis, which represents the most extreme form of *Lactobacillus* depletion and anaerobic overgrowth, is one of the strongest risk factors for preterm birth. And it explains why the risk is graded: women with *L. crispatus* dominance (maximum protection) have the lowest risk, women with *L. iners* dominance (intermediate protection) have moderate risk, and women with CST IV communities (minimal *Lactobacillus* protection) have the highest risk.

Fertility and the Endometrial Microbiome

The influence of the urogenital microbiome on reproduction begins before pregnancy. A growing body of evidence suggests that the microbial communities of the vagina, cervix, and endometrium influence fertility — both natural conception and assisted reproductive technologies (ART).

The concept of an **endometrial microbiome** is newer and more controversial than the vaginal microbiome. The endometrium — the lining of the uterus — was traditionally considered sterile, protected by the cervical mucus plug and the tight junction of the internal cervical os. Culture-independent studies have detected bacterial DNA in endometrial samples, with *Lactobacillus* and *Gardnerella* among the most commonly identified genera. However, whether these signals represent genuine endometrial colonisation or contamination from the cervix and vagina during sampling remains debated.

What is less debated is the clinical observation that the microbial environment of the reproductive tract influences implantation outcomes. In IVF (in vitro fertilisation) studies, women whose endometrial or vaginal microbiome is *Lactobacillus*-dominated at the time of embryo transfer have been reported to have higher implantation rates, higher clinical pregnancy rates, and lower miscarriage rates than women with dysbiotic (non-*Lactobacillus*) communities. These findings are suggestive rather than definitive — sample sizes are small, confounding factors are numerous, and the causal direction is unclear. But they fit the broader pattern: a *Lactobacillus*-dominated reproductive tract seems to provide a more hospitable environment for the establishment and maintenance of pregnancy.

The Prevention Question

If vaginal dysbiosis contributes to preterm birth, can restoring a *Lactobacillus*-dominated community prevent it? This is one of the most pressing questions in perinatal microbiome research — and the answer, so far, is frustratingly incomplete.

Screening and treatment of bacterial vaginosis during pregnancy has been tested in multiple randomised controlled trials, with mixed results. Some trials of antibiotic treatment for BV in early pregnancy showed modest reductions in preterm birth risk; others showed no benefit. The inconsistency likely reflects the heterogeneity of both BV (which is not a single condition but a spectrum of microbial states) and preterm birth (which has multiple aetiologies, only some of which are microbiome-related).

Probiotic interventions — administering *Lactobacillus* species directly to the vagina — have also been tested, with similarly inconsistent results. As we will explore in detail in Chapter 23, the challenges of probiotic colonisation are substantial: the introduced organisms must survive the vaginal environment, compete with the existing community, establish residence, and produce sufficient lactic acid to shift the pH. Many probiotic strains fail on one or more of these criteria.

The most promising avenue may be precision approaches: identifying women at high risk based on their vaginal microbiome profile (particularly those with CST IV communities or *L. iners* dominance in early pregnancy) and intervening with targeted therapies — whether antibiotics, probiotics, or live biotherapeutic products — tailored to their specific microbial ecology. But we are not there yet. The vaginal microbiome is a powerful predictor of preterm birth risk; translating that prediction into effective prevention remains the frontier.

8.4 Sex Differences in Microbiome Composition and Disease Risk

Throughout this book, we have described the human microbiome as if it were a single entity — “the” gut microbiome, “the” skin microbiome. But the microbiome is shaped by the host, and hosts come in (at least) two sexes. The question of whether and how the microbiome differs between men and women — and what consequences those differences have for health and disease — is one that microbiome science is only beginning to address.

The Obvious Differences

Some sex differences in the microbiome are straightforward consequences of anatomy. The vaginal microbiome, as we have just seen, is a uniquely female habitat. Men lack it entirely — and the penile microbiome, while far less studied, appears to be a more diverse, less *Lactobacillus*-dominant community shaped by circumcision status, sexual activity, and hygiene practices. The urobiome, too, likely differs between the sexes, though this has been less extensively characterised in men than in women.

The Subtler Differences — Gut and Beyond

More interesting, and more clinically relevant, are the sex differences in microbial communities that both sexes share — particularly the gut.

Multiple studies have reported differences in gut microbiome composition between men and women, though the findings are not always consistent. A pattern that emerges across several large cohorts is that premenopausal women tend to have higher relative abundances of certain *Lactobacillus* and *Bifidobacterium* species compared to age-matched men, and that these differences attenuate after menopause — suggesting, once again, a role for sex hormones in shaping the microbiome.

The mechanism linking sex hormones to the gut microbiome is bidirectional. Oestrogen and testosterone can influence microbial community composition directly (by altering the intestinal epithelium, mucus production, and immune function) and indirectly (through effects on diet, metabolism, and bile acid composition). But the microbiome also influences sex hormone levels — a concept that has been formalised as the **estrobolome**.

The estrobolome, proposed by Plottel and Blaser in 2011, refers to the aggregate of enteric bacterial genes whose products are capable of metabolising oestrogens. Certain gut bacteria produce **β -glucuronidase**, an enzyme that deconjugates oestrogen metabolites in the gut, allowing them to be reabsorbed into the circulation in their active form rather than being excreted. A gut microbiome with high β -glucuronidase activity therefore increases circulating oestrogen levels; a depleted or dysbiotic microbiome with low activity decreases them. This has potential implications for oestrogen-dependent conditions including breast cancer, endometriosis, and osteoporosis — though the clinical evidence remains largely associational [REF:plottel2011].

Sex Differences in Microbiome-Associated Disease

The most compelling evidence for sex-differentiated microbiome effects comes from disease states in which incidence or outcomes differ between the sexes and in which the microbiome has been implicated.

Autoimmune diseases show a striking female preponderance — overall, approximately 80 per cent of autoimmune disease patients are women. Conditions including rheumatoid arthritis, systemic lupus erythematosus, and multiple sclerosis are all far more common in women. There are exceptions - some autoimmune conditions such as myocarditis appear more common in men. The microbiome has been implicated in the pathogenesis of several of these conditions (a story we will explore in detail in Chapter 26), and some researchers have proposed that sex differences in gut microbial composition contribute to the female bias in autoimmune disease. In animal models, the gut microbiome has been shown to modulate the sex hormone-dependent development of autoimmune diabetes, and transfer of male microbiota to female mice can alter their hormone profiles and protect them from autoimmune disease — a phenomenon that Markle and colleagues termed the “sex-microbiome-autoimmunity” axis.

Cardiovascular disease, conversely, is more common in men — or more precisely, in premenopausal women it is less common than in age-matched men, with the gap narrowing after menopause. The gut microbial metabolite **trimethylamine N-oxide (TMAO)**, produced from dietary choline and carnitine by gut bacteria and subsequently oxidised in the liver, has been identified as a risk factor for atherosclerosis and cardiovascular events. Sex differences in TMAO levels have been reported, with some studies finding higher levels in men, potentially mediated by differences in gut microbial composition, diet, and the interaction between sex hormones and the flavin monooxygenase (FMO3) enzyme that converts TMA to TMAO [REF:razavi2019].

Colorectal cancer also shows sexual dimorphism — it is more common in men, and men tend to be diagnosed at earlier ages. The concept of a “microgenderome” — sex-specific interactions between the gut microbiome and the host — has been invoked to explain some of this disparity. Oestrogen, by promoting a more favourable gut microbiota (higher *Lactobacillus* and *Bifidobacterium*, lower proportion of certain gram-negative species), may confer a degree of protection; androgens, conversely, may promote an environment more conducive to the pro-carcinogenic species, including *Fusobacterium nucleatum*, that we will encounter in Chapter 28 [REF:wu2024].

The Gut-Brain-Microbiome Axis and Sex

Even the gut-brain axis — the bidirectional communication between the gut microbiome and the central nervous system (Chapter 13) — shows sex-dependent patterns. Jašarević, Morrison, and Bale have reviewed the evidence for sex differences in the gut microbiome-brain axis across the lifespan, noting that the microbiome undergoes sexually dimorphic development that parallels sex

differences in brain maturation. Stress exposure, which is a major modulator of gut microbial composition, has sex-specific effects on both the microbiome and on neuropsychiatric outcomes. The implications for understanding sex differences in depression, anxiety, and neurodevelopmental disorders are tantalising — though, as with much of this field, we are still in the phase of identifying associations rather than proving causation [REF:jasarevic2016].

A Cautionary Note

It is important to end this section with a caveat. The study of sex differences in the microbiome is in its infancy. Many studies are underpowered, fail to account for confounding variables (diet, medication use, BMI, age), or do not adequately control for the anatomical differences (the vaginal microbiome's spillover effect on the gut, for instance) that could drive spurious associations. The largest and most robust human microbiome studies to date — including the Human Microbiome Project and several large population-based cohorts — have generally found that sex explains a relatively small proportion of the total variation in gut microbiome composition compared to factors like diet, geography, medication use, and individual genetics.

This does not mean sex differences are unimportant. It means they are subtle, context-dependent, and likely mediated through indirect pathways — particularly hormonal — rather than through a single, dramatic mechanism. The lesson of the microbiome, as always, is ecological: the system is complex, the variables interact, and simplistic narratives (“men have microbiome X, women have microbiome Y”) are likely to be misleading.

Where This Matters: If you are a clinician managing conditions in which the microbiome has been implicated — whether it is recurrent urinary tract infections, irritable bowel syndrome, autoimmune disease, or cardiovascular risk — it is worth remembering that your male and female patients may not be starting from the same microbial baseline. Sex-specific reference ranges for the microbiome do not yet exist, but the evidence suggests they will eventually be needed — just as we now recognise that reference ranges for blood counts, liver enzymes, and cardiac biomarkers differ between the sexes. The microbiome is a biological system shaped by the body it inhabits, and that body is sexed.

In the next chapter, we explore a body site where the microbiome has been studied for longer than anywhere else in the body — the mouth. From dental plaque to periodontal disease to the newly appreciated oral-gut axis, the oral microbiome turns out to be not just the oldest field in microbiome research but one of the most clinically relevant.

Chapter 8 — Key References

- [REF:ravel2011] Ravel, J., Gajer, P., Abdo, Z., Schneider, G.M., Koenig, S.S.K., McCulle, S.L., ... & Forney, L.J. (2011). Vaginal microbiome of reproductive-age women. *Proceedings of the National Academy of Sciences*, 108(Suppl 1), 4680–4687. DOI: 10.1073/pnas.1002611107
- [REF:gajer2012] Gajer, P., Brotman, R.M., Bai, G., Sakamoto, J., Schütte, U.M.E., Zhong, X., ... & Ravel, J. (2012). Temporal dynamics of the human vaginal microbiota. *Science Translational Medicine*, 4(132), 132ra52. DOI: 10.1126/scitranslmed.3003605

- [REF:smith2017] Smith, S.B. & Ravel, J. (2017). The vaginal microbiota, host defence and reproductive physiology. *Journal of Physiology*, 595(2), 451–463. DOI: 10.1113/JP271694
- [REF:price2020] Price, T.K., Hilt, E.E., Thomas-White, K., Mueller, E.R., Wolfe, A.J. & Brubaker, L. (2020). The urobiome of continent adult women: a cross-sectional study. *BJOG*, 127(2), 193–201. DOI: 10.1111/1471-0528.15920
- [REF:wolfe2022] Joyce, C., Halverson, T., Gonzalez, C., Brubaker, L. & Wolfe, A.J. (2022). The urobiomes of adult women with various lower urinary tract symptoms status differ: a re-analysis. *Frontiers in Cellular and Infection Microbiology*, 12, 860408. DOI: 10.3389/fcimb.2022.860408
- [REF:thomas-white2020] Thomas-White, K., Taege, S., Limeira, R., Brincat, C., Joyce, C., Hilt, E.E., ... & Wolfe, A.J. (2020). Vaginal estrogen therapy is associated with increased Lactobacillus in the urine of postmenopausal women with overactive bladder symptoms. *American Journal of Obstetrics and Gynecology*, 223(5), 727.e1–727.e11. DOI: 10.1016/j.ajog.2020.08.006
- [REF:gudnadottir2022] Gudnadottir, U., Debelius, J.W., Du, J., Hugerth, L.W., Danielsson, H., Schuppe-Koistinen, I., Fransson, E. & Brusselaers, N. (2022). The vaginal microbiome and the risk of preterm birth: a systematic review and network meta-analysis. *Scientific Reports*, 12(1), 7926. DOI: 10.1038/s41598-022-12007-9
- [REF:huang2023] Huang, C., Gin, C., Fettweis, J., Foxman, B., Gelaye, B., MacIntyre, D.A., ... & Callahan, B. (2023). Meta-analysis reveals the vaginal microbiome is a better predictor of earlier than later preterm birth. *BMC Biology*, 21(1), 199. DOI: 10.1186/s12915-023-01702-2
- [REF:ihmp2019] The Integrative Human Microbiome Project Research Network Consortium. (2019). The Integrative Human Microbiome Project. *Nature*, 569(7758), 641–648. DOI: 10.1038/s41586-019-1238-8
- [REF:kenneally2022] Kenneally, C., Murphy, C.P., Sleator, R.D. & Culligan, E.P. (2022). The urinary microbiome and biological therapeutics: novel therapies for urinary tract infections. *Microbiological Research*, 259, 127010. DOI: 10.1016/j.micres.2022.127010
- [REF:plottel2011] Plottel, C.S. & Blaser, M.J. (2011). Microbiome and malignancy. *Cell Host & Microbe*, 10(4), 324–335. DOI: 10.1016/j.chom.2011.10.003
- [REF:razavi2019] Razavi, A.C., Potts, K.S., Kelly, T.N. & Bazzano, L.A. (2019). Sex, gut microbiome, and cardiovascular disease risk. *Biology of Sex Differences*, 10(1), 29. DOI: 10.1186/s13293-019-0240-z
- [REF:jasarevic2016] Jašarević, E., Morrison, K.E. & Bale, T.L. (2016). Sex differences in the gut microbiome-brain axis across the lifespan. *Philosophical Transactions of the Royal Society B*, 371(1688), 20150122. DOI: 10.1098/rstb.2015.0122
- [REF:wu2024] Wu, Z., Huang, Y., Zhang, R., Zheng, C., You, F., Wang, M., Xiao, C. & Li, X. (2024). Sex differences in colorectal cancer: with a focus on sex hormone-gut microbiome axis. *Cell Communication and Signaling*, 22(1), 167. DOI: 10.1186/s12964-024-01549-2

Additional references to be verified and expanded during manuscript preparation.

Chapter 9: The Oral Microbiome — More Than Just Teeth

In the previous chapter, we explored the urogenital tract — a body system where the assumption of sterility persisted for decades and where its dismantling has transformed our understanding of reproductive health. Now we turn to a body site where the story of microbiology itself began, and where the microbiome has been studied for longer than anywhere else in the human body.

On a September day in 1683, the Dutch draper and self-taught lens-grinder Antonie van Leeuwenhoek scraped a sample of white matter from between his own teeth, placed it under one of his hand-crafted microscopes, and became the first human being in history to observe living bacteria. He described what he saw in a letter to the Royal Society of London: “little living animalcules, very prettily a-moving.” Some rolled, some shot through the water “like a pike,” and others spun like a top. Leeuwenhoek had no germ theory, no concept of infection, and no way of knowing that the substance he had scraped from his teeth — dental plaque — would turn out to be one of the most complex and clinically consequential microbial ecosystems on the planet. But in that single observation, he had founded both microbiology and oral microbiome research simultaneously.

More than three centuries later, the mouth remains one of the best-characterised microbial habitats in the human body. The **Human Oral Microbiome Database (HOMD)**, established by Floyd Dewhirst, Tsute Chen, and colleagues at the Forsyth Institute, catalogues over 700 species-level taxa across 13 phyla — making the oral cavity the second most diverse microbial habitat in the body after the gut [REF:dewhirst2010]. But the oral microbiome is not merely diverse. It is ancient, architecturally sophisticated, and clinically significant in ways that extend far beyond cavities and gum disease. Oral bacteria influence cardiovascular health, contribute to gut inflammation, may play a role in neurodegeneration, and — because saliva is so easy to collect — offer a non-invasive diagnostic window into diseases affecting the entire body.

This chapter explores the oral microbiome in four parts. We begin with its most distinctive feature — the biofilm — and the dental plaque that was the original subject of microbiome research. We then examine periodontal disease and its surprising connections to cardiovascular and systemic health. From there, we follow the oral-gut axis: the daily swallowing of roughly a trillion bacteria and what happens when oral organisms colonise the intestine. Finally, we consider the emerging use of the oral microbiome as a diagnostic tool — a vision in which a sample of saliva might reveal diseases far removed from the mouth itself.

9.1 Biofilms and Dental Plaque: The Original Microbiome Research

A City, Not a Crowd

If you run your tongue across your teeth after a day without brushing, the slightly fuzzy film you feel is dental plaque — and it is not a random accumulation of bacteria. It is a **biofilm**: a highly organised, three-dimensional community of microorganisms embedded in a self-produced matrix of extracellular polymers. To call plaque a “film of germs” is like calling a city a “collection of people.” The description is technically accurate and entirely misleading.

Biofilms are the dominant mode of bacterial life on Earth. They coat the surfaces of rocks in rivers, the insides of water pipes, and the hulls of ships. In the human body, biofilms form on teeth, on implanted medical devices, and — as we saw in earlier chapters — on the mucosal surfaces of the gut and respiratory tract. But dental plaque holds a special place in biofilm science: it was the first microbial biofilm to be recognised as such, and it remains one of the most intensively studied.

What makes dental plaque a biofilm rather than simply a collection of bacteria sitting on a surface? Several features.

First, the community is **structured**. Bacteria in dental plaque are not randomly distributed. They are spatially organised, with specific species occupying specific positions relative to the tooth surface and to one another. Early colonisers — principally *Streptococcus* species and *Actinomyces* — attach directly to a proteinaceous conditioning film called the **acquired enamel pellicle** that coats the tooth surface within seconds of cleaning. These pioneers then provide attachment sites for later arrivals, building the community upward and outward through a process called **coaggregation**: the specific, receptor-mediated binding of one bacterial species to another, like Velcro strips that only stick to their matching partners [REF:valm2019].

Second, the community is **embedded in a matrix**. The bacteria produce an extracellular scaffold of polysaccharides, proteins, lipids, and extracellular DNA that holds the community together, provides mechanical stability, retains water and nutrients, and — critically — shields the bacteria within from both the host immune system and from antimicrobial agents. This matrix is one reason why antibiotics that kill planktonic (free-floating) bacteria in a test tube often fail to eliminate the same species when they are living in a biofilm. The matrix acts as a physical and chemical barrier, limiting drug penetration and creating microenvironments — gradients of oxygen, pH, and nutrients — that allow metabolically diverse species to coexist in close proximity [REF:zijnge2012].

Third, the community is **cooperative**. Different species within the biofilm perform different metabolic functions, and the waste products of one species often become the food of another. This metabolic cross-feeding creates a division of labour analogous to a functioning economy. *Streptococcus* species, for example, produce lactic acid as a byproduct of sugar fermentation. *Veillonella* species, which cannot ferment sugars themselves, consume that lactic acid as their primary carbon source. The result is a partnership: *Streptococcus* feeds *Veillonella*, and *Veillonella* removes lactic acid that would otherwise lower the pH and damage the tooth enamel. This is mutualism playing out at a micrometre scale [REF:valm2019].

The Architecture of Plaque

For decades, the internal architecture of dental plaque was poorly understood. Researchers could identify which species were present through culture or DNA sequencing, but the spatial relationships between species — who was next to whom, and how the community was physically organised — remained hidden.

That changed dramatically with the development of advanced fluorescence imaging techniques, particularly **combinatorial labelling and spectral imaging fluorescence in situ hybridisation (CLASI-FISH)**. This method, developed by Jessica Mark Welch, Gary Borisy, and colleagues at the Forsyth Institute and the Marine Biological Laboratory, allows researchers to label up to fifteen different bacterial taxa simultaneously with distinct fluorescent probes and then image their spatial arrangement in intact biofilm samples at single-cell resolution.

The images that emerged were breathtaking. Supragingival plaque — the plaque that forms on the

exposed surfaces of teeth above the gum line — turned out to be organised into elaborate, multi-taxon structures that the researchers described with evocative names. A “**hedgehog**” structure consists of a central core of *Corynebacterium* filaments radiating outward from the tooth surface, with distinct bacterial taxa arranged in concentric zones around them: *Streptococcus* species on the outer periphery, *Porphyromonas* and *Haemophilus* in intermediate zones, and *Lautropia* near the tips. The whole assembly — visible under the microscope as a radially organised sphere roughly 100 to 200 micrometres in diameter — resembles a miniature hedgehog or sea urchin.

Other structures include “**corncob**” formations, in which cocci (round bacteria, typically *Streptococcus* or *Actinomyces*) adhere in neat rows along the length of a central filamentous organism, producing a structure that looks exactly like a cob of maize under the microscope. These corncob formations had actually been observed in dental plaque as early as the 1970s, but their significance — as evidence of specific, reproducible inter-species physical organisation — was not fully appreciated until the CLASI-FISH era.

The message from these imaging studies is clear: dental plaque is not a disorganised smear of bacteria. It is a spatially structured, multi-species community in which physical proximity reflects functional relationships. Species that metabolically cooperate tend to be physically adjacent. Species that compete for the same resources tend to be spatially separated. The biofilm is, in a very real sense, a microbial tissue — with a level of organisation that we more typically associate with the cells of a multicellular organism.

The Ecological Niches of the Mouth

The mouth is not a single habitat. It is a landscape of strikingly different ecological niches, each harbouring its own distinct microbial community. Consider the diversity of surfaces and conditions within the oral cavity:

The **teeth** provide the only non-shedding hard surfaces in the body — surfaces on which biofilms can accumulate without being periodically removed by the turnover of epithelial cells. This is why plaque forms on teeth and not on the tongue or cheeks: the epithelial surfaces of the oral mucosa shed their outermost cells roughly every 12 to 24 hours, taking any attached bacteria with them. Tooth enamel does not shed. A biofilm established on a tooth surface can persist and mature indefinitely — unless mechanically removed by brushing, flossing, or professional cleaning.

The **tongue**, particularly its dorsal (upper) surface, is the largest reservoir of bacteria in the mouth. The papillae of the tongue — the small bumps that give it its rough texture — create a vast surface area of crypts and folds in which bacteria accumulate. The tongue dorsum is the primary source of the volatile sulphur compounds (hydrogen sulphide and methyl mercaptan) responsible for **halitosis** — bad breath. It is also, as we shall see, the principal site of the nitrate-reducing bacteria that play a surprising role in cardiovascular physiology.

The **gingival sulcus** — the shallow groove between each tooth and the gum tissue that surrounds it — is an environment radically different from the exposed tooth surface just millimetres away. The sulcus is bathed in **gingival crevicular fluid (GCF)**, a serum-like exudate that provides proteins, immunoglobulins, and complement factors. It is warmer, moister, and more anaerobic than the supragingival tooth surface. The bacterial community of the gingival sulcus is correspondingly different: more anaerobic, more gram-negative, and — when the sulcus deepens into a **periodontal pocket** — more dominated by the species associated with gum disease.

The **saliva** itself, rather than being a habitat, functions as a transport medium — a river flowing

through the oral landscape, carrying bacteria from one site to another, delivering antimicrobial proteins (lysozyme, lactoferrin, secretory IgA), buffering pH changes, and providing the fluid in which nutrients and metabolic products are exchanged. A healthy adult produces roughly 0.5 to 1.5 litres of saliva per day. Each millilitre contains approximately 100 million bacterial cells. The saliva does not so much host a community as sample from every community in the mouth — which, as we shall see in section 9.4, is precisely what makes it so interesting as a diagnostic fluid.

The hard and soft **palate**, the **buccal mucosa** (inner cheeks), and the **tonsils** each harbour distinct communities that differ from one another and from the teeth, tongue, and gingival sulcus. The tonsils, with their deep crypts, are particularly rich in anaerobic organisms and are a significant source of oral bacterial diversity.

The Dewhirst HOMD catalogue identified over 700 species-level taxa across these sites, of which only about a quarter have been formally named and cultured. Another 8 per cent have been cultured but remain unnamed, and a remarkable 68 per cent are known only from their DNA sequences — detected by 16S rRNA gene sequencing but never grown in a laboratory [REF:dewhirst2010]. The oral microbiome, despite being one of the longest-studied microbial communities on Earth, still contains a majority of organisms that no one has ever succeeded in growing.

The Socransky Complexes — A Map of Gum Disease

In 1998, Sigmund Socransky and Anne Haffajee, working at the Forsyth Dental Center in Boston, published a landmark study that would reshape the field of periodontal microbiology. Using DNA-DNA hybridisation to analyse over 13,000 subgingival plaque samples from 185 subjects, they identified five distinct clusters — or **complexes** — of bacterial species that consistently co-occurred in subgingival plaque. They colour-coded these complexes, and the colour scheme has become one of the most widely recognised classification systems in oral microbiology [REF:socransky1998].

The **yellow complex** (*Streptococcus* species), **green complex** (*Capnocytophaga*, *Eikenella*, *Campylobacter concisus*), and **purple complex** (*Veillonella*, *Actinomyces*) are associated with health and early plaque formation. They are the pioneers — the first colonisers of a clean tooth surface.

The **orange complex** — including *Fusobacterium nucleatum*, *Prevotella intermedia*, *Prevotella nigrescens*, and *Peptostreptococcus micros* — represents a bridging community. These species are associated with the transition from health to disease. *Fusobacterium nucleatum* plays a particularly important role as a physical and metabolic bridge: it coaggregates with both the early colonisers and the late colonisers, physically linking the health-associated bacteria to the disease-associated ones. Without *Fusobacterium*, the pathogenic species struggle to integrate into the biofilm.

The **red complex** — *Porphyromonas gingivalis*, *Tannerella forsythia*, and *Treponema denticola* — is the group most strongly associated with severe periodontal disease. These three species were consistently found together, and their presence correlated strikingly with deeper pocket depths, greater clinical attachment loss, and more bleeding on probing. They are the late arrivals — organisms that require the anaerobic, nutrient-rich environment created by the preceding colonisers — and they are the organisms most closely linked to the tissue destruction of advanced periodontitis [REF:socransky1998].

The Socransky model has been refined and complicated by subsequent molecular studies — the relationships between species are more fluid and context-dependent than a simple colour scheme

suggests, and organisms outside the original five complexes are now recognised as important contributors to disease. But the fundamental insight remains valid: periodontal disease is not caused by a single pathogen. It is an ecological shift — a **dysbiosis** — in which the entire community structure changes, driven by positive feedback loops between microbial metabolism and host inflammation. Understanding this shift is the subject of our next section.

Where This Matters: Every time you brush your teeth, you are disrupting a biofilm that took roughly 24 hours to mature — a structured, spatially organised, metabolically cooperative community that was rebuilding itself before you had even put the toothbrush down. You are not sterilising your mouth. You cannot, and you should not. What you are doing is preventing the community from reaching the mature, complex state in which the late colonisers — the red complex organisms associated with periodontal disease — can establish themselves. Brushing does not eliminate bacteria. It resets the ecological clock.

9.2 Periodontal Disease and Systemic Health — The Cardiovascular Connection

When the Ecosystem Breaks Down

Periodontal disease — disease of the gums and the supporting structures of the teeth — is one of the most prevalent chronic diseases in the world. In its milder form, **gingivitis**, the gums become inflamed, reddened, and prone to bleeding. In its more severe form, **periodontitis**, the inflammation extends deeper, destroying the periodontal ligament and the alveolar bone that anchor the teeth in the jaw. Untreated, periodontitis leads to tooth loss. Globally, severe periodontitis affects roughly 10 to 15 per cent of adults, and some form of periodontal disease is estimated to affect nearly half the adult population [REF:kinane2017].

The pathogenesis of periodontal disease illustrates the ecological principles we have been building throughout this book. Periodontitis is not, in the classical sense, an infection — not in the way that tuberculosis is caused by *Mycobacterium tuberculosis* or cholera by *Vibrio cholerae*. There is no single exogenous pathogen that invades from outside and causes disease. The bacteria responsible for periodontitis are already there. They are members of the commensal oral microbiome — organisms that coexist with their host in health and that cause tissue destruction only when the ecological balance of the community shifts.

The shift begins with plaque accumulation. When oral hygiene is neglected — when the biofilm is allowed to mature undisturbed for days or weeks — the community structure changes. The early colonisers (streptococci, actinomyces) are joined and eventually outnumbered by the later-arriving anaerobes. The subgingival environment becomes more anaerobic, more protein-rich (as the gingival crevicular fluid flow increases in response to inflammation), and more dominated by the orange and red complex organisms. *Porphyromonas gingivalis*, *Tannerella forsythia*, and *Treponema denticola* flourish.

But these organisms do not destroy tissue directly by producing toxins in the way a classic pathogen might. Instead, they provoke and manipulate the host immune response. *Porphyromonas gingivalis* is a particularly sophisticated operator. It produces proteases called **gingipains** that degrade host proteins, disrupt complement pathways, and manipulate Toll-like receptor signalling. It can invade epithelial cells and persist intracellularly, evading immune clearance. And it exerts

what researchers have called a “**keystone pathogen**” effect — an influence on the community that is disproportionate to its abundance. Even at low levels, *P. gingivalis* can alter the host immune response in ways that destabilise the entire microbial community, shifting it from a health-compatible state to a dysbiotic, disease-promoting one [REF:kinane2017].

The resulting inflammation is the primary driver of tissue damage. Neutrophils and macrophages are recruited to the gingival sulcus. Pro-inflammatory cytokines — interleukin-1 β , interleukin-6, tumour necrosis factor- α — are released. Matrix metalloproteinases (MMPs) degrade the collagen of the periodontal ligament. Osteoclasts are activated, resorbing the alveolar bone. The gingival sulcus deepens into a periodontal pocket — a self-reinforcing process, because the deeper pocket creates a more anaerobic environment that further favours the pathogenic anaerobes. This is the positive feedback loop — the “vicious cycle” — of periodontitis: dysbiosis drives inflammation, inflammation reshapes the local environment, and the reshaped environment drives further dysbiosis [REF:rosier2017].

The disease is modulated by risk factors that will be familiar from earlier chapters: smoking (the single strongest modifiable risk factor for periodontitis, roughly tripling the risk), diabetes (particularly poorly controlled diabetes, which impairs neutrophil function and increases gingival inflammation), genetic susceptibility, stress, and the composition of the diet. But the fundamental mechanism is ecological: a breakdown in the homeostatic relationship between the biofilm community and the host immune response.

Beyond the Mouth — The Systemic Connection

If periodontal disease were merely a dental problem — painful and inconvenient, but confined to the mouth — it would still be one of the most important microbial diseases in the world by virtue of its prevalence alone. But over the past two decades, a growing body of evidence has linked periodontitis to a range of systemic conditions, most prominently **cardiovascular disease**.

The association between poor oral health and heart disease was first noted in epidemiological studies in the late 1980s and 1990s. Since then, it has been confirmed in numerous large cohort studies and meta-analyses. A 2020 meta-analysis synthesising data from prospective studies found that individuals with periodontitis had a roughly 20 to 30 per cent increased risk of cardiovascular events (myocardial infarction, stroke, cardiovascular death) compared to those without periodontitis, even after adjustment for shared risk factors such as smoking, diabetes, age, and socioeconomic status [REF:demmer2006].

The association is clear. The critical question — and one that remains actively debated — is whether the relationship is **causal**. Three broad mechanisms have been proposed.

First: bacteraemia. The inflamed periodontal pocket is essentially an open wound bathed in a dense bacterial biofilm. Everyday activities — chewing, brushing teeth, even flossing — can introduce oral bacteria directly into the bloodstream. This transient **bacteraemia** is well documented and occurs more frequently and at higher levels in individuals with periodontitis than in those with healthy gums. Once in the bloodstream, oral bacteria can reach distant sites. And they have been found there. DNA and viable cells of *Porphyromonas gingivalis*, *Fusobacterium nucleatum*, *Aggregatibacter actinomycetemcomitans*, and other oral species have been detected in **atherosclerotic plaques** — the fatty deposits in arterial walls that underlie coronary artery disease and stroke. Their presence in atherosclerotic lesions has been confirmed by multiple independent groups using both PCR-based detection of bacterial DNA and, in some studies, culture of viable organisms from

plaque tissue [REF:demmer2006].

This is a striking finding. The bacteria that Socransky colour-coded as the “red complex” of gum disease are turning up inside the arterial walls of patients with heart disease. The implication — that oral pathogens may directly contribute to the development or destabilisation of atherosclerotic plaques — remains an active area of investigation.

Second: systemic inflammation. Periodontitis is a chronic inflammatory condition that elevates systemic markers of inflammation. C-reactive protein (CRP), interleukin-6 (IL-6), and fibrinogen — all of which are established risk factors for cardiovascular disease — are consistently elevated in patients with periodontitis compared to periodontally healthy controls. The hypothesis is that the chronic, low-grade inflammation originating in the periodontal tissues contributes to the systemic inflammatory burden that drives atherosclerosis. This does not require bacteria to physically travel to the arterial wall. The inflammatory mediators themselves — carried through the bloodstream from the inflamed gums — may be sufficient to promote endothelial dysfunction, accelerate lipid deposition, and destabilise existing plaques.

Third: shared risk factors and confounding. Sceptics have argued that the association between periodontitis and cardiovascular disease may be explained, at least in part, by shared risk factors that predispose to both conditions. Smoking, diabetes, obesity, socioeconomic deprivation, and psychological stress all increase the risk of both periodontitis and cardiovascular disease. If these confounders are not fully controlled for in epidemiological studies, the apparent link between oral and cardiovascular disease could be spurious — a statistical artefact rather than a biological reality.

The honest assessment, as of this writing, is that the evidence supports a genuine biological association between periodontitis and cardiovascular disease that goes beyond simple confounding — but falls short of proving that treating periodontitis reduces cardiovascular risk. Intervention studies — trials in which patients with periodontitis receive intensive periodontal treatment and are then followed for cardiovascular outcomes — have generally shown improvements in surrogate markers of cardiovascular risk (reductions in CRP, improvements in endothelial function) but have not yet demonstrated definitive reductions in hard endpoints such as heart attacks or strokes. The studies that have been conducted are relatively small and of limited duration. Larger, longer trials are underway, but the question remains open [REF:demmer2006].

Periodontal Disease and the Nitrate Paradox

There is a second, less well-known cardiovascular connection — and it runs in the opposite direction.

The tongue dorsum, as noted in section 9.1, is home to a consortium of nitrate-reducing bacteria, including species of *Veillonella*, *Actinomyces*, *Haemophilus*, *Neisseria*, and *Rothia*. These organisms play a key role in the **enterosalivary nitrate cycle** — a physiological pathway with direct implications for blood pressure regulation.

The cycle works as follows. Dietary nitrate, abundant in green leafy vegetables and beetroot, is absorbed from the gut into the bloodstream and concentrated in the salivary glands, which secrete it back into the mouth at concentrations 10 to 20 times higher than those in plasma. On the tongue dorsum, nitrate-reducing bacteria convert the salivary nitrate to **nitrite**. This nitrite is swallowed with saliva and, in the acidic environment of the stomach and in peripheral tissues, is further reduced to **nitric oxide (NO)** — a potent vasodilator. The net effect of this cycle is to supplement

the body's endogenous nitric oxide production, contributing to the maintenance of healthy blood pressure [REF:pignatelli2020].

The clinical evidence for this pathway is compelling. Studies have shown that the use of antiseptic mouthwashes — particularly those containing **chlorhexidine** — can disrupt the oral nitrate-reducing bacterial community and, within days, produce measurable increases in blood pressure in otherwise healthy individuals. Conversely, dietary nitrate supplementation (the “beetroot juice” studies) lowers blood pressure through a mechanism that is at least partly dependent on oral bacterial nitrate reduction. Kill the tongue bacteria with mouthwash, and the blood pressure-lowering effect of dietary nitrate is abolished [REF:morouBermudez2022].

This creates an intriguing paradox for periodontal disease. Periodontitis is associated with a dysbiotic shift in the oral microbiome that may impair the nitrate-reducing community. At the same time, aggressive antiseptic mouthwash use — a common recommendation for managing periodontal disease — may further deplete these beneficial bacteria. The net result could be a reduction in nitric oxide bioavailability and a consequent increase in cardiovascular risk — adding a second, microbiome-mediated pathway to the already complex relationship between oral and cardiovascular health [REF:pignatelli2020].

Where This Matters: The next time you reach for an antiseptic mouthwash, it is worth pausing to consider what you are killing. Most over-the-counter mouthwashes are indiscriminate — they reduce the total bacterial load of the mouth without distinguishing between the organisms that promote gum disease and those that contribute to blood pressure regulation. For most healthy individuals, mechanical cleaning — brushing and flossing — is sufficient to disrupt the plaque biofilm and prevent periodontal disease without collateral damage to the nitrate-reducing community. The advice “rinse, don’t think” may need revision.

9.3 The Oral-Gut Axis: Swallowing a Trillion Microbes a Day

The Great Migration

Every time you swallow, you send a bolus of saliva — and with it, somewhere between 100 million and a billion bacteria — on a one-way journey from the mouth to the stomach. You swallow roughly 600 to 1,000 times per day. The arithmetic is sobering: over the course of 24 hours, you deliver an estimated **1.5 trillion** oral bacteria to your gastrointestinal tract. This is not an occasional event. It is a continuous, lifelong microbial migration — a river of organisms flowing from the most accessible microbial habitat in the body to the largest.

For decades, this migration was assumed to be biologically inconsequential. The reasoning was straightforward: the stomach, with its pH of 1.5 to 3.5, should kill the vast majority of swallowed bacteria. And the few survivors that reached the small intestine and colon would face a well-established resident community that, through the mechanism of **colonisation resistance** we discussed in Chapter 5, should prevent newcomers from gaining a foothold. Oral bacteria, the thinking went, are adapted to the mouth. They lack the metabolic machinery to thrive in the anaerobic, bile-rich, nutrient-competitive environment of the gut. They pass through as transients and are eliminated.

This view has turned out to be substantially wrong — or at least, substantially incomplete.

Oral Bacteria in the Gut — More Than Transients

The first clue that oral bacteria might be doing more than passing through came from metagenomic studies of the gut microbiome that noticed something unexpected: a significant fraction of the species detected in the gut — particularly in patients with certain diseases — were species previously known only from the oral cavity.

A landmark study by Koji Atarashi, Kenya Honda, and colleagues, published in *Science* in 2017, provided the most compelling experimental evidence to date. Working with gnotobiotic (germ-free) mice, the researchers isolated *Klebsiella* strains from the saliva of a healthy human volunteer and introduced them into the gut. These oral *Klebsiella* strains — which are ordinarily minor components of the oral microbiota — colonised the mouse gut readily and proved to be potent inducers of **TH1 (T helper 1) immune cells**, a lineage of immune cells associated with inflammation. When these same strains were introduced into genetically susceptible mice, they triggered severe intestinal inflammation resembling inflammatory bowel disease [REF:atarashi2017].

Crucially, the oral *Klebsiella* strains could only colonise the gut when the resident intestinal microbiome was disrupted — either by antibiotic treatment or by the germ-free condition of the gnotobiotic mice. In mice with an intact, healthy gut microbiome, the oral bacteria were excluded by colonisation resistance. But when that resistance was compromised — when the ecological defences of the gut were down — the oral invaders established themselves and caused disease.

This finding reframes the daily swallowing of oral bacteria in a new light. Under normal conditions, the healthy gut microbiome acts as a barrier, preventing oral species from colonising the intestine even though they arrive there continuously. But under conditions of dysbiosis — following antibiotic treatment, during inflammatory bowel disease, or in other states in which the gut community is perturbed — oral bacteria can breach this barrier, take up residence in the gut, and drive inflammation. The mouth, as Atarashi and colleagues put it, serves as a **reservoir of potential intestinal pathobionts** — organisms that are harmless in their native habitat but pathogenic when they colonise the wrong one.

Fusobacterium nucleatum — The Quintessential Oral-Gut Traveller

No species better illustrates the oral-gut axis than *Fusobacterium nucleatum*. This organism, which we encountered in section 9.1 as the “bridge” species in the Socransky complexes — the coaggregation hub that links early and late colonisers in the dental biofilm — has emerged as one of the most clinically significant bacteria in the human body, and its significance extends far beyond the mouth.

F. nucleatum is a normal member of the oral microbiome. In the mouth, it performs an ecological role: it physically and metabolically connects different members of the biofilm community, facilitating the development of complex plaque. But *F. nucleatum* has also been detected, with remarkable consistency, in **colorectal tumours**. Study after study has found *F. nucleatum* enriched in colorectal cancer tissue compared to adjacent normal tissue — a finding that has been replicated across dozens of independent cohorts worldwide.

The mechanistic evidence is increasingly detailed. In cell culture and animal models, *F. nucleatum* promotes colorectal cancer through multiple pathways: it adheres to and invades intestinal epithelial cells via the FadA adhesin and Fap2 lectin; it activates the β -catenin/Wnt signalling pathway, promoting cell proliferation; it modulates the tumour immune microenvironment by recruiting myeloid-derived suppressor cells and inhibiting natural killer cell and T cell activity; and it

promotes chemoresistance by activating autophagy pathways. These findings have elevated *F. nucleatum* from a bystander to a potential driver — or at least a facilitator — of colorectal carcinogenesis. We will return to this story in detail in Chapter 28 [REF:lam2023].

The question, for our purposes here, is: how does an oral bacterium end up in a colorectal tumour? Two routes have been proposed. The first is the **haematogenous route** — the same bacteraemia mechanism we discussed in the context of cardiovascular disease. *F. nucleatum* enters the bloodstream through inflamed periodontal tissues and is carried to the colon, where it adheres to tumour tissue via the Fap2 protein, which binds to the sugar Gal-GalNAc that is overexpressed on colorectal cancer cells. The second route is **enteral** — the organism is swallowed, survives the gastric acid barrier (perhaps sheltered within a bolus of food or within biofilm aggregates), and reaches the colon via the oral-gut axis [REF:lam2023].

Neither route has been definitively proven in humans, and it is possible that both operate simultaneously. But the clinical correlation is robust: patients with periodontal disease have higher levels of *F. nucleatum* in the oral cavity and, in several studies, higher levels in colorectal tumour tissue. The mouth-to-gut journey of *Fusobacterium nucleatum* — from dental biofilm bridge organism to colorectal cancer accomplice — is one of the most dramatic examples of how the oral microbiome can influence health far beyond the oral cavity.

***Porphyromonas gingivalis* — From Gums to Gut to Brain?**

The story of *Porphyromonas gingivalis* extends even further. We have already discussed its role as a keystone pathogen in periodontitis and its detection in atherosclerotic plaques. But in 2019, a study by Stephen Dominy, Casey Lynch, Jan Potempa, and colleagues made international headlines with a finding that pushed the boundaries of oral-systemic connection: *P. gingivalis* and its gingipain proteases were detected in the **brain tissue of patients with Alzheimer's disease** [REF:dominy2019].

The study, published in *Science Advances*, reported that gingipains — the signature virulence factors of *P. gingivalis* — were present in more than 90 per cent of the Alzheimer's brain samples examined, and that their levels correlated with the burden of tau pathology and ubiquitin pathology. In a mouse model, oral infection with *P. gingivalis* led to brain colonisation by the bacterium and increased production of amyloid- β , the protein that forms the characteristic plaques of Alzheimer's disease. Small-molecule inhibitors of gingipains reduced the bacterial load in the brain, blocked amyloid- β production, and rescued neurons in the hippocampus [REF:dominy2019].

These findings must be treated with appropriate caution. The study was partly funded by CorTexyme, a pharmaceutical company developing gingipain inhibitors for Alzheimer's treatment — the drug ultimately failed in clinical trials. And the question of whether *P. gingivalis* is a cause, a contributor, or merely a bystander in Alzheimer's disease remains unresolved. The organism could be opportunistically colonising brain tissue already damaged by the disease rather than initiating the damage. Correlation, as we have emphasised throughout this book, is not causation.

But the broader pattern is worth noting. *P. gingivalis* has now been detected in atherosclerotic plaques, in brain tissue, in rheumatoid joints, and in colorectal tumour tissue. Whether it is causing disease at these distant sites or merely taking advantage of damaged tissue, the fact that an oral bacterium turns up so consistently in diseased tissue throughout the body underlines a fundamental point: the mouth is not a closed system. The oral microbiome is connected to every other organ in the body — by the bloodstream, by the gut, and by the immune system.

The Oral-Gut Axis in Inflammatory Bowel Disease

The connection between oral bacteria and intestinal inflammation is particularly well documented in **inflammatory bowel disease (IBD)**. Epidemiological studies have consistently found a higher prevalence of periodontal disease in patients with IBD compared to healthy controls — and conversely, a higher prevalence of IBD in patients with periodontitis. The relationship appears to be bidirectional: periodontal disease may worsen gut inflammation, and gut inflammation may alter the oral microbiome [REF:zhou2023].

The mechanistic link runs through the ectopic colonisation pathway demonstrated by Atarashi. In patients with Crohn's disease and ulcerative colitis, the gut microbiome is already dysbiotic — colonisation resistance is compromised. Oral bacteria, particularly *P. gingivalis*, *F. nucleatum*, and *Klebsiella* species, can exploit this vulnerability to colonise the gut. Once there, they activate TH1 and TH17 immune responses, amplifying the intestinal inflammation that characterises IBD. The result is a potential vicious cycle: gut dysbiosis permits oral bacterial colonisation, oral bacteria drive inflammation, inflammation worsens dysbiosis, and the cycle continues [REF:lam2023] [REF:zhou2023].

This oral-gut axis also connects to the mycobiome thread we have been following throughout the book. *Candida* species, which are normal residents of the oral cavity, can similarly translocate to the gut and expand during periods of bacterial dysbiosis — particularly after antibiotic use. The interkingdom dynamics of the oral-gut axis — bacteria and fungi migrating together from mouth to gut, competing and cooperating in their new environment — is an area of active research that we will revisit in Chapter 10.

Where This Matters: The idea that what happens in your mouth stays in your mouth is a fiction. Every day, your mouth seeds your gut with trillions of bacteria. Most of the time, your gut community repels these immigrants. But when the gut's defences are down — after antibiotics, during inflammatory disease, or in states of immune compromise — oral bacteria can colonise the intestine and drive or exacerbate disease. This connection has practical implications: maintaining good oral health is not just about saving your teeth. It is about reducing the reservoir of potentially pathogenic organisms that your gut must contend with every day.

9.4 The Oral Microbiome as a Diagnostic Window

Why Saliva?

Medicine has long depended on blood. A blood draw — venepuncture — is the standard gateway to clinical diagnostics: blood counts, liver function, kidney function, lipid profiles, inflammatory markers, tumour markers, genetic tests. The procedure is safe, quick, and well standardised. It is also invasive, requires trained personnel, causes discomfort, and is difficult to perform repeatedly for longitudinal monitoring. For home-based, point-of-care, or resource-limited settings, blood-based diagnostics are often impractical.

Saliva, by contrast, is the most accessible biofluid in the human body. It can be collected non-invasively, painlessly, and repeatedly by anyone — no needles, no trained phlebotomist, no special equipment. A patient can spit into a tube at home and mail the sample to a laboratory. A child can provide a sample without distress. A public health programme can screen thousands of people with

minimal infrastructure.

But accessibility means nothing without information content. The critical question is: does saliva contain enough biological information to be diagnostically useful?

The answer, increasingly, is yes — and the reason lies in the multiple pathways by which molecules from the blood and the rest of the body find their way into saliva.

Saliva is not simply water and enzymes. It is a filtrate of the blood. The major salivary glands (parotid, submandibular, sublingual) and the hundreds of minor salivary glands scattered across the oral mucosa are richly vascularised. Molecules from the bloodstream — proteins, nucleic acids, metabolites, hormones, antibodies — cross from the capillaries into the salivary gland acinar cells and from there into the saliva through a combination of passive diffusion, active transport, and ultrafiltration. In addition, gingival crevicular fluid — the serum-like exudate from the gingival sulcus — flows continuously into the oral cavity, contributing plasma-derived molecules directly to the saliva.

The result is that saliva contains a remarkably comprehensive molecular profile. The salivary **proteome** includes over 3,000 proteins, many of which are shared with the plasma proteome. The salivary **transcriptome** includes mRNA and microRNA species that reflect gene expression in both local oral tissues and distant organs. The salivary **metabolome** captures hundreds of small-molecule metabolites. And the salivary **microbiome** — the sum of all microbial DNA present in the saliva — samples from every microbial community in the oral cavity. Collectively, these molecular layers have been given the umbrella term “**salivaomics**” — a term coined by David Wong and colleagues at UCLA, who have been among the most active proponents of saliva as a diagnostic medium [REF:wong2012].

Salivary Diagnostics for Oral Disease

The most straightforward application of salivary diagnostics is in the detection of oral diseases themselves. Since saliva bathes the oral cavity, it directly reflects the local microbial and inflammatory environment.

For **dental caries**, salivary levels of *Streptococcus mutans* and *Lactobacillus* species — the principal acidogenic organisms in the caries biofilm — have been used as risk indicators for decades, though their predictive value at the individual level is modest. More recent metagenomic and metabolomic approaches have identified broader microbial signatures and metabolic profiles (elevated lactic acid, reduced ammonia production) that distinguish caries-active from caries-free individuals with better accuracy [REF:rosier2017].

For **periodontal disease**, salivary biomarkers hold greater promise. Salivary levels of pro-inflammatory cytokines (IL-1 β , IL-6, TNF- α), matrix metalloproteinases (MMP-8, MMP-9), and tissue degradation products have been shown to correlate with periodontal disease severity. The detection of specific periodontal pathogens — *P. gingivalis*, *T. forsythia*, *T. denticola* — in saliva can confirm the presence of a dysbiotic subgingival microbiome. And the gingipains of *P. gingivalis*, whose role in tissue destruction we discussed in section 9.2, are detectable in saliva and gingival crevicular fluid and have been proposed as targets for point-of-care diagnostic devices.

For **oral cancer**, salivary diagnostics are perhaps most urgently needed. Oral squamous cell carcinoma is frequently diagnosed at an advanced stage, when treatment options are limited and five-year survival rates are poor. Salivary biomarkers — including specific mRNA transcripts, microR-

NAs, and proteins — have been identified that can distinguish oral cancer patients from healthy controls in case-control studies. The challenge, as with all cancer biomarkers, is achieving sufficient sensitivity and specificity for population-level screening. But the non-invasive nature of saliva collection makes it an attractive medium for the development of screening tools, particularly for high-risk populations (tobacco users, heavy alcohol consumers, individuals with HPV-related oropharyngeal risk) [REF:martina2020].

Saliva as a Window to Systemic Disease

The more transformative — and more challenging — vision of salivary diagnostics extends beyond the mouth entirely. If saliva contains molecules derived from the bloodstream, can it be used to detect diseases affecting distant organs?

The evidence is accumulating across several disease areas.

Cardiovascular disease. Given the periodontal-cardiovascular connection discussed in section 9.2, it is perhaps unsurprising that salivary markers of inflammation and vascular risk have been investigated. Salivary CRP, salivary IL-6, and salivary levels of cardiac biomarkers (such as cardiac troponin and brain natriuretic peptide) have been detected, though at much lower concentrations than in blood and with substantial analytical challenges. The clinical utility of salivary cardiovascular biomarkers remains to be demonstrated in large prospective studies.

Diabetes. Salivary glucose levels do not reliably correlate with blood glucose levels — the salivary glands do not simply filter glucose from the plasma. However, salivary markers of oxidative stress, inflammatory cytokines, and — intriguingly — the composition of the salivary microbiome itself have been associated with type 2 diabetes risk. The observation that chronic antiseptic mouthwash use is associated with increased risk of prediabetes, presumably through disruption of the nitrate-reducing oral bacterial community, adds another dimension to the oral-metabolic connection [REF:morouBermudez2022].

Cancer beyond the oral cavity. Perhaps the most ambitious application of salivary diagnostics is the detection of cancers in distant organs. Wong and colleagues at UCLA have reported the identification of salivary transcriptomic biomarkers that can distinguish patients with pancreatic cancer and lung cancer from healthy controls. The mechanistic basis for the appearance of tumour-specific mRNA in saliva is not fully understood — it may involve exosomes (small membrane-bound vesicles) released by tumour cells into the bloodstream that are subsequently taken up by salivary glands and released into saliva [REF:wong2012].

Infectious disease. The COVID-19 pandemic dramatically accelerated interest in saliva as a diagnostic medium. Salivary SARS-CoV-2 testing — detecting viral RNA in saliva rather than in nasopharyngeal swabs — proved to be a viable alternative to swab-based PCR, with comparable sensitivity and the enormous practical advantage of self-collection. Saliva-based testing for other respiratory and systemic infections, including tuberculosis and viral hepatitis, is an active area of development [REF:khurshid2020].

Neuropsychiatric conditions. Salivary cortisol, a well-established marker of hypothalamic-pituitary-adrenal (HPA) axis activity, has been used for decades in stress research and in the diagnosis of Cushing's syndrome. More recently, salivary levels of alpha-synuclein — a protein implicated in Parkinson's disease — and of certain microRNAs associated with neurodegeneration have been investigated as potential early markers of neurodegenerative disease. These applications are

early-stage and unvalidated for clinical use, but they illustrate the breadth of the diagnostic ambitions for saliva.

Challenges and the Road Ahead

For all its promise, salivary diagnostics faces significant challenges that have slowed its translation from research to clinical practice.

The first is **standardisation**. Saliva composition varies with the time of day, the time since the last meal, hydration status, salivary flow rate, the method of collection (passive drool versus stimulated saliva), and the specific glands contributing to the sample. Without standardised collection and processing protocols, biomarker measurements can vary widely between studies and between samples from the same individual.

The second is **sensitivity**. Most blood-derived biomarkers are present in saliva at concentrations 100 to 1,000 times lower than in blood. Detecting them requires highly sensitive analytical platforms — mass spectrometry, digital PCR, electrochemical biosensors — that are not yet available in most clinical laboratories and are far from the point-of-care simplicity that makes saliva attractive as a diagnostic medium in the first place.

The third is **specificity**. The oral cavity is not a passive conduit. It is a metabolically active environment in which oral bacteria produce, modify, and degrade molecules. A protein detected in saliva might originate from the blood, from oral tissues, or from the oral microbiome. Distinguishing between these sources — and between disease-specific signals and the normal background variation of the oral ecosystem — is a non-trivial analytical challenge [REF:martina2020].

Despite these challenges, the trajectory is clear. Advances in omics technologies, biosensor engineering, and artificial intelligence are steadily closing the gap between the information content of saliva and our ability to extract it. Point-of-care salivary diagnostic devices — capable of measuring multiple biomarkers simultaneously from a single drop of saliva and returning results within minutes — are in advanced stages of development. When they arrive, they will transform the mouth from the oldest subject of microbiome research into one of the newest frontiers of precision diagnostics.

Where This Matters: The mouth has always been the easiest part of the body to examine. You open it, and a clinician can see teeth, gums, tongue, and mucosa without any equipment more sophisticated than a light and a mirror. Saliva extends this accessibility to the molecular level. If the current trajectory of research continues, a sample of your saliva — collected painlessly in a few seconds — may one day provide information about your cardiovascular risk, your metabolic health, and the early stages of diseases that have not yet produced symptoms. Your dentist's office may become a screening station for conditions that currently require a blood draw and a visit to a specialist.

In the next chapter, we turn to the forgotten kingdom — the fungi. Throughout our tour of the body's microbial habitats, we have spoken almost exclusively of bacteria. But at every site we have visited — the gut, the skin, the respiratory tract, the urogenital tract, and now the mouth — fungi have been present, interacting with bacterial communities and with the host immune system in ways we are only beginning to understand. Chapter 10 reframes everything we have learned so far: go back through those sites, and realise that the mycobiome was there all along.

Chapter 9 — Key References

- [REF:dewhirst2010] Dewhirst, F.E., Chen, T., Izard, J., Paster, B.J., Tanner, A.C.R., Yu, W.-H., Lakshmanan, A. & Wade, W.G. (2010). The human oral microbiome. *Journal of Bacteriology*, 192(19), 5002–5017. DOI: 10.1128/JB.00542-10
- [REF:socransky1998] Socransky, S.S., Haffajee, A.D., Cugini, M.A., Smith, C. & Kent, R.L. (1998). Microbial complexes in subgingival plaque. *Journal of Clinical Periodontology*, 25(2), 134–144. DOI: 10.1111/j.1600-051x.1998.tb02419.x
- [REF:kinane2017] Kinane, D.F., Stathopoulou, P.G. & Papapanou, P.N. (2017). Periodontal diseases. *Nature Reviews Disease Primers*, 3, 17038. DOI: 10.1038/nrdp.2017.38
- [REF:valm2019] Valm, A.M. (2019). The structure of dental plaque microbial communities in the transition from health to dental caries and periodontal disease. *Journal of Molecular Biology*, 431(16), 2957–2969. DOI: 10.1016/j.jmb.2019.05.016
- [REF:zijnge2012] Zijnge, V., Ammann, T., Thurnheer, T. & Gmür, R. (2012). Subgingival biofilm structure. *Frontiers of Oral Biology*, 15, 1–16. DOI: 10.1159/000329667
- [REF:rosier2017] Rosier, B.T., Marsh, P.D. & Mira, A. (2017). Resilience of the oral microbiota in health: mechanisms that prevent dysbiosis. *Journal of Dental Research*, 97(4), 371–380. DOI: 10.1177/0022034517742139
- [REF:demmer2006] Demmer, R.T. & Desvarieux, M. (2006). Periodontal infections and cardiovascular disease: the heart of the matter. *Journal of the American Dental Association*, 137(Suppl), 14S–20S. DOI: 10.14219/jada.archive.2006.0402
- [REF:pignatelli2020] Pignatelli, P., Fabietti, G., Ricci, A., Piattelli, A. & Curia, M.C. (2020). How periodontal disease and presence of nitric oxide reducing oral bacteria can affect blood pressure. *International Journal of Molecular Sciences*, 21(20), 7538. DOI: 10.3390/ijms21207538
- [REF:morouBermudez2022] Morou-Bermúdez, E., Torres-Colón, J.E., Bermúdez, N.S., Patel, R.P. & Joshipura, K.J. (2022). Pathways linking oral bacteria, nitric oxide metabolism, and health. *Journal of Dental Research*, 101(6), 623–631. DOI: 10.1177/00220345211064571
- [REF:atarashi2017] Atarashi, K., Suda, W., Luo, C., Kawaguchi, T., Motoo, I., Narushima, S., ... & Honda, K. (2017). Ectopic colonization of oral bacteria in the intestine drives TH1 cell induction and inflammation. *Science*, 358(6361), 359–365. DOI: 10.1126/science.aan4526
- [REF:lam2023] Lam, G.A., Albarrak, H., McColl, C.J., Pizarro, A., Sanaka, H., Gomez-Nguyen, A., Cominelli, F. & Paes Batista da Silva, A. (2023). The oral-gut axis: periodontal diseases and gastrointestinal disorders. *Inflammatory Bowel Diseases*, 29(7), 1153–1164. DOI: 10.1093/ibd/izac241
- [REF:zhou2023] Zhou, T., Xu, W., Wang, Q., Jiang, C., Li, H., Chao, Y., Sun, Y. & A, L. (2023). The effect of the “oral-gut” axis on periodontitis in inflammatory bowel disease: a review of microbe and immune mechanism associations. *Frontiers in Cellular and Infection Microbiology*, 13, 1132420. DOI: 10.3389/fcimb.2023.1132420
- [REF:dominy2019] Dominy, S.S., Lynch, C., Ermini, F., Benedyk, M., Marczyk, A., Konradi, A., ... & Potempa, J. (2019). *Porphyromonas gingivalis* in Alzheimer’s disease brains: evid-

ence for disease causation and treatment with small-molecule inhibitors. *Science Advances*, 5(1), eaau3333. DOI: 10.1126/sciadv.aau3333

- [REF:wong2012] Wong, D.T.W. (2012). Salivaomics. *Journal of the American Dental Association*, 143(10 Suppl), 19S–24S. DOI: 10.14219/jada.archive.2012.0339
- [REF:martina2020] Martina, E., Campanati, A., Diotallevi, F. & Offidani, A. (2020). Saliva and oral diseases. *Journal of Clinical Medicine*, 9(2), 466. DOI: 10.3390/jcm9020466
- [REF:khurshid2020] Khurshid, Z., Warsi, I., Moin, S.F., Slowey, P.D., Latif, M., Zohaib, S. & Zafar, M.S. (2020). Biochemical analysis of oral fluids for disease detection. *Advances in Clinical Chemistry*, 100, 205–253. DOI: 10.1016/bs.acc.2020.04.005
- [REF:deng2025] Deng, L., Guan, G., Cannon, R.D. & Mei, L. (2025). Age-related oral microbiota dysbiosis and systemic diseases. *Microbial Pathogenesis*, 205, 107717. DOI: 10.1016/j.micpath.2025.107717

Additional references to be verified and expanded during manuscript preparation.

Chapter 10: The Mycobiome — The Forgotten Kingdom

For nine chapters, we have toured the microbial landscape of the human body — the gut, the skin, the lungs, the urogenital tract, the mouth — and spoken almost exclusively of bacteria. We have described vast bacterial communities, traced their metabolic contributions, and marvelled at the sophistication of their interactions with the immune system. Bacteria have dominated the story because they dominate the research. But they are not the only residents.

At every site we have visited, fungi were there too. They were there when we described the skin, where the lipophilic yeast *Malassezia* outnumbers every other microbe on most of the body surface. They were there in the gut, where *Candida* species ride alongside the great bacterial majority. They were there in the lungs, where *Aspergillus* spores are inhaled with every breath. They were there in the mouth, where *Candida albicans* colonises the mucosal surfaces in roughly half the healthy population. And they were there in the urogenital tract, where vaginal candidiasis is one of the most common infections women experience.

Yet fungi have been, until very recently, almost invisible in microbiome science. This chapter asks why — and then sets out to correct the oversight. We will meet the major fungal inhabitants of the human body, explore how they interact with their bacterial neighbours and with the immune system, and examine what happens when the fungal community is disrupted. The story that emerges is one of a kingdom that has been underestimated at every turn: small in numbers, enormous in influence, and central to the immune education that keeps us healthy.

The term for this fungal community — the **mycobiome** — was itself only coined in 2010, when Iliyan Iliev, David Underhill, and colleagues at Cedars-Sinai Medical Center in Los Angeles published the first study demonstrating that gut fungi interact with the innate immune system through the receptor Dectin-1, and that disruption of this interaction worsens intestinal inflammation [REF:iliev2012]. That paper, published in *Science*, did not merely describe a fungal community. It named it, characterised its immunological footprint, and — in a single set of experiments — demonstrated that ignoring fungi meant ignoring a critical piece of the inflammatory puzzle. The mycobiome had arrived.

10.1 Why Fungi Were Overlooked: Technical Challenges in Sequencing and Culture

The Wrong Net

If you want to know what lives in the ocean, the answer depends on what kind of net you use. A net with a fine mesh catches small fish and plankton. A net with a coarse mesh lets them slip through. For decades, the tools used to study the human microbiome were, metaphorically speaking, designed to catch bacteria — and they let the fungi slip through.

The revolution in microbiome research was built on a single gene: the **16S ribosomal RNA gene**. As we discussed in Chapter 4, this gene is present in all bacteria and archaea, it contains regions that are conserved enough to serve as universal primer binding sites and variable enough to distinguish

species, and it can be amplified by PCR and sequenced to produce a census of every bacterial species in a sample. The 16S approach was elegant, powerful, and enormously productive. It generated the data that launched the Human Microbiome Project and transformed our understanding of the human body's microbial inhabitants.

But 16S is a bacterial gene. Fungi do not have it. The equivalent marker for fungi is a different DNA region: the **internal transcribed spacer (ITS)**, a stretch of non-coding DNA situated between the genes encoding the small and large ribosomal RNA subunits. The ITS region serves a similar purpose to 16S — it is universal across fungi, variable enough to permit species-level identification, and amplifiable by PCR — but it requires different primers, different PCR conditions, and different bioinformatic pipelines for analysis. A study designed to profile bacteria using 16S primers will detect zero fungi. The fungi are there in the sample. The method simply cannot see them.

This is not a subtle point. It means that the vast majority of early microbiome studies — the landmark papers that established which microbes live where, the large-scale projects that catalogued the “healthy” human microbiome — were blind to fungi from the start. Not because fungi were absent, but because nobody thought to look. Or more precisely: looking for fungi required a separate, parallel effort with different tools, different expertise, and different funding — and for years, that effort was simply not made.

The methodological gap ran deeper than primer choice. The reference databases for fungal identification were — and in many respects remain — far less complete than those for bacteria. The SILVA and Greengenes databases for bacterial 16S sequences contain millions of entries and are well curated. The UNITE database for fungal ITS sequences, while a remarkable community effort, has historically been smaller, less complete, and more prone to taxonomic ambiguity. Many environmental fungi have been sequenced but not formally named. Many clinical isolates have ITS sequences that match no known reference. The result is that even when researchers did perform ITS sequencing, they often obtained a higher proportion of unclassifiable sequences than they would in a bacterial 16S study — leading to the frustrating situation in which you know something is there, but you cannot say what it is [REF:zhang2022].

Culture-based methods — growing fungi in the laboratory from patient samples — faced their own challenges. Many gut fungi are strict anaerobes or require specialised media that were not part of standard clinical microbiology protocols. *Candida* species grow readily on Sabouraud dextrose agar, but many other gut-associated fungi do not. The culturomics revolution that has expanded the catalogue of cultivable bacteria (discussed in Chapter 4) has only recently been extended to fungi, and the fungal culture catalogue remains far smaller than its bacterial counterpart.

The cumulative effect of these technical biases was profound. For the first decade of the microbiome era — roughly 2005 to 2015 — fungi were effectively absent from the scientific conversation. Papers described “the microbiome” and meant “the bacteriome.” Review articles discussed “microbial diversity” and counted only bacteria. Clinical studies linked “dysbiosis” to disease and measured only bacterial community shifts. The mycobiome existed, but it existed in the dark.

Emerging from the Shadows

The situation began to change around 2013, with a handful of landmark studies that applied ITS sequencing to human samples for the first time at scale. The most influential was a study by Keisha Findley, Julia Segre, Heidi Kong, and colleagues at the National Human Genome Research Institute, who mapped the fungal communities across 14 skin sites in 10 healthy adults — the first com-

prehensive topographic survey of the human skin mycobiome [REF:findley2013]. Their findings, published in *Nature*, revealed a striking pattern: core body and arm sites were overwhelmingly dominated by a single genus, *Malassezia*, with species-level variation between sites. But foot sites — the plantar heel, toenail, and toe web — harboured a far more diverse fungal community, including *Aspergillus*, *Cryptococcus*, *Rhodotorula*, *Epicoccum*, and many others. The paper made clear that the skin mycobiome was not a uniform blanket of *Malassezia* but a patchwork of distinct communities shaped by local physiology — much as the bacterial skin microbiome had proven to be.

In the gut, the picture was initially more confused. Early ITS studies of faecal samples reported widely varying fungal compositions, with inconsistencies between cohorts and considerable debate about which fungi were genuine residents and which were transient passengers arriving with food. *Saccharomyces cerevisiae* — baker's yeast — turns up reliably in gut mycobiome surveys, but it is also present in bread, beer, and fermented foods. Is it a gut resident or a dietary visitor? The answer, as a comprehensive review by Fen Zhang, Tao Zuo, and colleagues in *The Lancet Microbe* made clear, is probably both: some *Saccharomyces* strains are genuinely resident, while others are transient [REF:zhang2022]. The distinction matters because the composition of the gut mycobiome is influenced not only by the body's internal ecology but by what you eat — a variable that is far more difficult to control for in human studies than most researchers initially appreciated.

10.2 Who's There? The Major Players

Despite the methodological challenges, a decade of ITS sequencing, metagenomics, and targeted culture studies has now produced a reasonably consistent picture of which fungi inhabit the human body. The catalogue is far shorter than the bacterial catalogue — the gut mycobiome comprises perhaps 50 to 100 regularly detected species, compared with over a thousand bacterial species — but the inhabitants are no less interesting.

Candida — The Ubiquitous Commensal

Candida is the genus most commonly associated with human fungal disease, but it is first and foremost a commensal — an organism that lives on and in the human body without causing harm under normal circumstances. *Candida albicans*, the most prevalent species, colonises the oral cavity, gastrointestinal tract, and vaginal mucosa of roughly 40 to 60 per cent of healthy adults. It is typically the most abundant fungal species detected in the gut by ITS sequencing and is the dominant commensal fungus of the human body across all age groups [REF:zhang2022].

As we described in Primer Chapter 4, *C. albicans* is **dimorphic** — it can switch between a rounded yeast form and an invasive hyphal form. In health, the yeast form predominates: small, budding cells that coexist peacefully with the mucosal epithelium and the bacterial microbiome. The switch to the hyphal form — triggered by changes in pH, temperature, nutrient availability, or host immune status — is associated with tissue invasion and disease. This morphological plasticity is central to *Candida*'s dual identity as commensal and pathogen, and we will return to it repeatedly in this chapter.

Other *Candida* species found in the human mycobiome include *Candida tropicalis*, *Candida parapsilosis*, *Candida glabrata* (now reclassified as *Nakaseomyces glabratus*), and *Candida dubliniensis*. Each has its own ecological niche, its own relationship with the immune system, and its own

clinical significance. *C. tropicalis*, for instance, has emerged as a species particularly relevant to inflammatory bowel disease — its abundance in the gut is increased in patients with Crohn's disease, and it was one of the species highlighted in the original Iliev et al. study as interacting with Dectin-1 [REF:iliev2012].

***Saccharomyces* — The Baker, the Brewer, and the Probiotic?**

Saccharomyces cerevisiae and its close relative *Saccharomyces boulardii* are among the most commonly detected fungi in the gut. As noted above, the distinction between genuine gut resident and dietary transient is not always clear for *Saccharomyces*. What is clear is that *S. boulardii* has been marketed extensively as a probiotic yeast, with claimed benefits for preventing antibiotic-associated diarrhoea and treating acute gastroenteritis. We will examine the evidence for these claims in Chapter 23. For now, we note that the presence of *Saccharomyces* in the gut is near-universal in Western populations and that it interacts with both the bacterial microbiome and the immune system in ways that are only beginning to be characterised.

A clinically important footnote: the antibodies that the immune system produces against *Saccharomyces* — known as **anti-*Saccharomyces cerevisiae* antibodies (ASCA)** — are used diagnostically as a biomarker for Crohn's disease. The fact that antibodies against a gut fungus serve as a marker for an inflammatory bowel disease is itself a powerful indication that the mycobiome is immunologically relevant. We will return to ASCA in section 10.7.

***Malassezia* — The Skin Specialist**

Malassezia is the dominant fungal genus on human skin, accounting for more than 90 per cent of the skin mycobiome at most body sites. These lipophilic yeasts thrive on the sebaceous (oil-rich) areas of the body — the scalp, face, chest, and back — where they feed on the lipids in sebum. As Findley and colleagues demonstrated, *Malassezia* essentially monopolises the fungal ecology of core body skin, with species including *Malassezia restricta*, *Malassezia globosa*, *Malassezia sympodialis*, and *Malassezia furfur* distributed across different anatomical sites [REF:findley2013].

Malassezia is also, however, a significant cause of skin disease — a duality that mirrors the commensal-pathogen spectrum we have seen in bacteria throughout this book. *Malassezia* species are implicated in **seborrheic dermatitis** (the chronic, flaking skin condition that in its mildest form is known as dandruff), **pityriasis versicolor** (the patchy skin discolouration common in tropical climates), and ***Malassezia* folliculitis** (inflammation of hair follicles). The mechanism is not straightforward overgrowth in the way that, say, a bacterial infection might be. Rather, *Malassezia* lipases break down skin lipids into free fatty acids, some of which — particularly oleic acid — trigger an inflammatory response in susceptible individuals. The same organism performing the same metabolic activity can be commensal in one person and pathogenic in another, depending on the host's immune and inflammatory response to its metabolic byproducts [REF:jung2023].

Intriguingly, *Malassezia* has also been detected in the gut — particularly in the gut mucosa of patients with Crohn's disease. A 2019 study by Limon and colleagues found *Malassezia restricta* enriched in the intestinal mucosa of Crohn's disease patients carrying the IBD-associated CARD9 risk allele, and demonstrated that *M. restricta* worsened colitis in a mouse model. This was a startling finding: a skin-adapted yeast, turning up in the inflamed gut and making things worse. We will explore this gut-skin-immune triangle in section 10.7.

***Aspergillus* — The Airborne Coloniser**

Aspergillus species are ubiquitous environmental moulds whose spores are inhaled with virtually every breath. A healthy person inhales hundreds of *Aspergillus* spores daily, the vast majority of which are cleared by the innate immune system without incident. But in immunocompromised individuals — particularly those with neutropenia, organ transplants, or advanced HIV — *Aspergillus fumigatus* can cause **invasive pulmonary aspergillosis**, a life-threatening infection with mortality rates exceeding 50 per cent even with treatment.

In the context of the mycobiome, *Aspergillus* occupies an interesting position. It is not a commensal in the traditional sense — it does not stably colonise healthy mucosal surfaces the way *Candida* or *Malassezia* do. Rather, it is a constant environmental challenger, a fungus that the immune system must continuously neutralise. The respiratory mycobiome, such as it is, is dominated by *Aspergillus* and other environmental moulds — *Cladosporium*, *Penicillium*, *Alternaria* — most of which are transient rather than resident. But their presence matters because the immune system's response to them — the constant, low-level antifungal surveillance — shapes the immunological tone of the respiratory tract. As we noted in Chapter 7, when this surveillance fails — as it did during the COVID-19 pandemic, when corticosteroid-treated patients developed devastating mucormycosis infections — the consequences are catastrophic.

10.3 Scale and Proportion: A Minority by Numbers, Not by Influence

One of the reasons fungi have been overlooked is that, by sheer numerical abundance, they are a small minority of the human microbiome. Estimates vary, but fungi are thought to represent roughly 0.1 per cent of the total microbial community in the gut by cell count. In a gut harbouring trillions of bacterial cells, the fungal population numbers in the billions — a large number in absolute terms, but dwarfed by its bacterial neighbours.

This numerical disparity has sometimes been taken to imply that fungi are unimportant — bit players on a bacterial stage. This inference is profoundly mistaken, for several reasons.

First, fungal cells are much larger than bacterial cells. A typical *Candida* yeast cell is 4 to 6 micrometres in diameter; a typical gut bacterium is 1 to 2 micrometres. The volume difference is roughly 10- to 50-fold. When fungi shift to hyphal growth, a single organism can extend filaments that are hundreds of micrometres long, penetrating tissue and crossing epithelial barriers. The physical presence of fungi is disproportionate to their numbers.

Second, fungi produce a vast array of bioactive metabolites — enzymes, toxins, signalling molecules — that influence the bacterial community, the host immune system, and the gut epithelium. A single fungal cell can exert biochemical effects on its surroundings that are out of all proportion to its contribution to total cell counts. The analogy is one of keystone species in ecology: just as a small population of wolves can reshape the behaviour of an entire elk herd and, through cascading effects, transform the vegetation of a river valley, a small population of fungi can reshape the microbial community of the gut.

Third — and most importantly — fungi are powerful stimulators of the immune system. The molecules that the immune system uses to detect fungi — including **β-glucan** in the fungal cell wall, recognised by the receptor **Dectin-1**, and **mannan** on the fungal surface, recognised by **Dectin-2** and the **mannose receptor** — are among the most potent activators of innate immunity known.

A small number of fungal cells can trigger an immune response that is far larger than what the same number of bacterial cells would provoke. This immunological visibility is not accidental. It reflects a billion years of co-evolution between fungi and the animal immune system, during which the immune system has learned to treat fungi with particular vigilance — a vigilance that, as we shall see, has profound consequences for both health and disease.

Where This Matters: The next time you read a headline about the “human microbiome,” notice what is being measured. If the study used 16S rRNA sequencing — as the majority of microbiome studies still do — it measured bacteria and nothing else. The fungi, the viruses, the archaea, the protists — the rest of the microbial community — were invisible. The mycobiome represents perhaps 0.1 per cent of gut microbes by cell count, but its immunological influence is vastly disproportionate to that number. Size, in microbiology, is not everything.

10.4 Fungal-Bacterial Interactions: Competition, Cooperation, and Cross-Kingdom Signalling

Fungi and bacteria do not exist in separate worlds within the human body. They cohabit the same mucosal surfaces, compete for the same nutrients, and communicate through chemical signals that cross the boundary between kingdoms. Understanding the mycobiome requires understanding these **interkingdom interactions** — the ways in which fungi and bacteria shape each other’s behaviour, abundance, and pathogenic potential.

Competition: The Bacterial Wall

The most fundamental interaction between gut bacteria and gut fungi is competition. A healthy, diverse bacterial microbiome acts as a powerful barrier against fungal overgrowth — a principle that is most dramatically illustrated by what happens when that barrier is removed.

When broad-spectrum antibiotics kill large swathes of the gut bacterial community, fungi — which are unaffected by antibacterial drugs — expand to fill the vacated ecological space. This phenomenon has been documented repeatedly in both animal models and human studies. The pioneering work of Mairi Noverr and Gary Huffnagle at the University of Michigan, published in *Infection and Immunity* in 2004, was among the first to demonstrate the principle experimentally. They showed that antibiotic treatment in mice led to stable increases in gut *Candida* levels, and that this fungal overgrowth — combined with subsequent mould spore exposure — could drive the development of allergic airway disease. The study provided the first experimental evidence that antibiotic disruption of gut bacteria could promote fungal expansion with consequences far beyond the gut itself [REF:noverr2004].

The mechanism of bacterial suppression of fungi involves several layers. Bacteria compete directly for nutrients and attachment sites. Certain bacterial species produce **antifungal metabolites** — short-chain fatty acids, particularly butyrate and propionate, inhibit *Candida* growth and hyphal formation. Lactic acid bacteria, including *Lactobacillus* species, produce lactic acid that lowers the local pH and creates conditions unfavourable for many fungi. And the bacterial community collectively maintains the integrity of the gut epithelial barrier, denying fungi access to the tissue beneath.

One particularly elegant example of cross-kingdom inhibition involves the bacterium *Enterococcus faecalis* and its secreted peptide **EntV**. Recent research has shown that EntV specifically inhibits the yeast-to-hypha transition in *Candida albicans* — the very morphological switch that is central to *Candida* pathogenicity. In effect, a commensal bacterium produces a molecule that keeps a commensal fungus in its harmless yeast form, preventing it from becoming invasive. This is inter-kingdom policing: the bacterial community actively suppresses fungal virulence [REF:zhang2022].

Cooperation: Strange Bedfellows

Not all fungal-bacterial interactions are antagonistic. In some contexts, fungi and bacteria cooperate — often to the detriment of the host.

Biofilm formation is a prime example. As we discussed in Chapter 9, biofilms are structured, surface-attached microbial communities encased in a protective matrix. In the oral cavity, *Candida albicans* can co-aggregate with oral streptococci to form mixed-species biofilms on dental prostheses and mucosal surfaces. The bacterial presence enhances *Candida* biofilm formation, and the *Candida* presence enhances bacterial persistence. These mixed biofilms are more resistant to antimicrobial treatment than either single-species community alone — a clinical problem familiar to anyone who has dealt with denture stomatitis or catheter-associated infections.

In the gut, the relationship between *Candida* and certain bacterial species can be similarly synergistic. During intestinal inflammation, the breakdown of the mucosal barrier creates opportunities for both fungi and bacteria to access tissue that is normally off-limits. The inflammatory milieu — rich in nutrients, low in oxygen, populated by activated immune cells — can favour the growth of organisms from both kingdoms that thrive in disturbed environments. The result is a mixed-kingdom dysbiosis in which fungi and bacteria amplify each other's effects on the host.

Cross-Kingdom Signalling

Perhaps the most fascinating aspect of fungal-bacterial interactions is the chemical dialogue that occurs between them. Both fungi and bacteria produce and detect small signalling molecules, and some of these molecules cross the kingdom boundary.

Quorum-sensing molecules — chemical signals that bacteria use to coordinate behaviour based on population density — can influence fungal behaviour. The *Pseudomonas aeruginosa* quorum-sensing molecule **3-oxo-C12-homoserine lactone** inhibits *Candida albicans* hyphal formation, keeping the fungus in its less virulent yeast form. Conversely, *Candida* produces **farnesol**, a quorum-sensing molecule of its own, which inhibits the production of certain virulence factors in *Pseudomonas*. The two organisms are, in effect, engaged in a chemical negotiation — each suppressing the other's most aggressive behaviours through reciprocal signalling.

These interactions are not confined to the laboratory. In the lungs of cystic fibrosis patients, *Pseudomonas aeruginosa* and *Candida albicans* co-exist in the airway mucus, and their interaction — competitive, cooperative, and communicative in turn — influences the course of pulmonary disease. In the gut, the signalling between bacterial and fungal communities is even more complex, involving multiple species on both sides and a host immune system that is simultaneously monitoring and responding to signals from both kingdoms.

The emerging picture is one of the human microbiome as a genuinely multi-kingdom ecosystem — not a collection of bacteria with a few fungi sprinkled on top, but an integrated community in which cross-kingdom interactions shape the behaviour of every member.

10.5 The Mycobiome and the Immune System: Dectin-1, CARD9, and Anti-Fungal Immunity as a Tonic Signal

Seeing Fungi: The Pattern Recognition Problem

The immune system faces a unique challenge with fungi. As we discussed in Primer Chapter 4, fungi are eukaryotes — their cells share much of the same basic molecular machinery as human cells. The immune system cannot rely on the same recognition strategies it uses for bacteria, whose prokaryotic cell walls (built from peptidoglycan) are utterly foreign to the eukaryotic body. Fungi have cell walls too, but they are built from different materials — chitin, β -glucan, and mannan — that require dedicated recognition systems.

The result is a specialised arm of innate immunity devoted to fungal detection, centred on a family of receptors called **C-type lectin receptors (CLRs)**. The most important of these for gut immunity is **Dectin-1** (also known as CLEC7A), a receptor on the surface of macrophages, dendritic cells, and neutrophils that recognises **β -1,3-glucan** — a structural polysaccharide found in the cell walls of virtually all fungi but absent from mammalian cells.

When Dectin-1 binds β -glucan, it triggers a signalling cascade that activates the transcription factor NF- κ B and induces the production of pro-inflammatory cytokines — including TNF- α , IL-6, and IL-1 β — as well as anti-microbial reactive oxygen species. This response is essential for controlling fungal infections. But Dectin-1 signalling is not just a defence mechanism triggered by fungal invasion. It is also activated by the commensal fungi that reside in the gut under normal conditions — providing a continuous, low-level immune stimulus that shapes the baseline activity of the intestinal immune system.

The CARD9 Hub

The signal from Dectin-1 — and from several other CLRs that detect fungal components — converges on an intracellular adaptor protein called **CARD9** (caspase activation and recruitment domain-containing protein 9). CARD9 acts as a signalling hub: it receives inputs from multiple fungal recognition receptors and translates them into an integrated immune response.

The importance of CARD9 was demonstrated dramatically by the discovery that humans with loss-of-function mutations in the *CARD9* gene are severely susceptible to fungal infections — particularly invasive candidiasis and dermatophytosis (fungal skin infections). These individuals cannot mount an effective immune response to fungi, even though their immunity to bacteria and viruses is largely intact. CARD9 deficiency is, in effect, a selective immunodeficiency for fungi — proof that the immune system has evolved a dedicated signalling pathway for this kingdom.

In 2021, a landmark study by Itai Doron, Irina Leonardi, Iliyan Iliev, and colleagues at Weill Cornell Medicine extended the CARD9 story into the mycobiome [REF:doron2021]. Using a technique they called **multi-kingdom antibody profiling (multiKAP)**, they mapped the human antibody repertoire against gut commensal fungi. They found that *Candida albicans* is the dominant inducer of systemic anti-fungal IgG antibodies — antibodies produced not in response to infection, but in response to normal gut colonisation. Fungal colonisation of the gut triggered germinal centre reactions in lymph nodes outside the intestine, generating systemic antibodies that conferred protection against disseminated *Candida* and *Candida auris* infection. Critically, this protective antibody

response depended on CARD9 — patients with CARD9 loss-of-function mutations had impaired anti-fungal IgG responses and were more vulnerable to invasive candidiasis.

This finding reframes the gut mycobiome as an immunological training ground. The commensal fungi in your gut are not passive bystanders. They are continuously stimulating your immune system through Dectin-1 and CARD9, driving the production of anti-fungal antibodies that protect you against potentially lethal fungal infections. The mycobiome is, in this sense, a **tonic signal** — a constant, low-level immune stimulus that keeps the anti-fungal arm of your immune system primed and ready to respond. Remove the signal — by eliminating gut fungi, or by crippling the CARD9 pathway — and the immune system becomes dangerously naive to fungal threats.

The Dectin-1 Polymorphism and Ulcerative Colitis

The original Iliev et al. study that gave the mycobiome its name made a second, equally important discovery. In addition to showing that mice lacking Dectin-1 were more susceptible to chemically induced colitis — because they could not properly regulate their response to commensal gut fungi — the researchers identified a polymorphism in the human *CLEC7A* gene (which encodes Dectin-1) that was associated with a severe form of ulcerative colitis [REF:iliev2012].

The implication was clear: the ability to recognise and respond to gut fungi is not merely a defence against fungal infection. It is a prerequisite for maintaining intestinal immune homeostasis. When the Dectin-1 pathway is impaired — whether by genetic polymorphism or experimental knockout — the immune system's relationship with the mycobiome breaks down, and the result is not fungal invasion but intestinal inflammation. The fungi themselves may not be pathogenic, but the immune system's failure to manage them properly is.

This insight — that anti-fungal immunity is not just about fighting infection but about maintaining the delicate balance between tolerance and response — is one of the most important conceptual contributions of mycobiome research to immunology.

Where This Matters: The immune system does not merely tolerate the fungi in your gut. It actively monitors them, using dedicated receptors (Dectin-1, Dectin-2, the mannose receptor) and a dedicated signalling pathway (CARD9) to maintain a continuous state of anti-fungal readiness. This is not paranoia — it is training. The commensal fungi in your gut prime your immune system to produce antibodies and immune cells that protect you against dangerous fungal infections. People with genetic defects in this pathway suffer devastating fungal diseases. The mycobiome is, quietly, one of the most important teachers your immune system has.

10.6 *Candida* Expansion During Antibiotic Use — The Opportunist's Playbook

Every physician knows the scenario. A patient receives a course of broad-spectrum antibiotics for a bacterial infection — pneumonia, a urinary tract infection, a wound infection. The antibiotics work: the bacterial infection clears. But a few days later, the patient develops a new problem. White patches on the tongue. An itchy vaginal discharge. Or, in hospitalised and immunocompromised patients, something far more sinister: *Candida* in the bloodstream.

This sequence — antibiotics eliminate bacteria, fungi fill the gap — is one of the clearest demonstrations of interkingdom ecology in the human body. It is also one of the oldest and most consistently

observed clinical phenomena in infectious disease medicine. And yet, for decades, it was treated as an unfortunate side effect rather than as an ecological principle with deep implications for how we think about the microbiome.

The Ecological Mechanism

The mechanism is straightforward in principle, though intricate in detail. A healthy gut bacterial community suppresses *Candida* through multiple overlapping mechanisms: nutrient competition, production of antifungal metabolites (particularly short-chain fatty acids), maintenance of an acidic pH, support of the epithelial barrier, and stimulation of immune pathways that hold *Candida* in its commensal yeast form.

Broad-spectrum antibiotics remove these bacterial competitors wholesale. The ecological niches previously occupied by bacteria — the attachment sites, the nutrient sources, the metabolic partnerships — are suddenly vacant. *Candida*, resistant to antibacterial drugs, expands to fill the void. In the absence of bacterial competition and the immune-modulating signals that bacteria provide, *Candida* may also undergo the yeast-to-hypha transition, switching from harmless commensal to tissue-invasive pathogen.

The Noverr and Huffnagle studies in the early 2000s demonstrated that this is not merely a gut-level phenomenon. In their mouse model, antibiotic-induced gut bacterial disruption combined with *Candida* colonisation led to an exaggerated immune response to subsequent airway allergen challenge — driving allergic airway disease through an IL-13-dependent mechanism [REF:noverr2005]. The gut mycobiome, when disrupted, could influence immune responses in the lungs. This was an early and prescient demonstration of the gut-lung axis — a connection that would become devastatingly relevant during the COVID-19 pandemic, when antibiotic- and corticosteroid-treated patients developed mucormycosis and invasive aspergillosis at alarming rates.

More recent work has revealed that the consequences of antibiotic-induced fungal expansion extend beyond the gut-lung axis. A 2025 study published in *The ISME Journal* showed that antibiotic-induced gut fungal overgrowth — predominantly *Candida* species — worsened the outcome of *Klebsiella pneumoniae* lung infections in mice through a Dectin-1-dependent mechanism. The expanded gut fungi upregulated Dectin-1 expression on alveolar macrophages, triggering excessive IL-1 β secretion and neutrophil recruitment that damaged lung tissue. Antifungal treatment or Dectin-1 knockout reversed the effect [REF:he2025]. The finding illuminated a disturbing feedback loop: antibiotics intended to treat bacterial pneumonia could, by promoting gut fungal overgrowth, worsen the very lung inflammation they were meant to resolve.

Clinical Implications

The clinical implications of antibiotic-induced fungal expansion are already well recognised in medicine, if not always framed in ecological terms. Oral and vaginal candidiasis are among the most common complications of antibiotic therapy. In intensive care units, *Candida* bloodstream infections (candidaemia) are a leading cause of nosocomial infection, with mortality rates of 30 to 50 per cent — and the majority of cases occur in patients who have received broad-spectrum antibiotics.

The ecological framing, however, adds something that the purely clinical perspective misses: the recognition that the problem is not merely *Candida* growth, but the loss of the bacterial community that normally keeps *Candida* in check. This reframing has therapeutic implications. If the fundamental problem is bacterial depletion, then the solution — or at least part of the solution — may

lie in restoring bacterial community function rather than simply adding antifungal drugs. This is the logic behind the use of *Saccharomyces boulardii* as a probiotic during antibiotic therapy — an approach we will evaluate critically in Chapter 23 — and behind broader efforts to develop microbiome-sparing antibiotic strategies that kill the target pathogen while preserving the commensal community that holds fungi in check.

Where This Matters: If you have ever taken a course of antibiotics and developed a yeast infection — oral thrush, vaginal candidiasis — you have experienced interkingdom ecology first-hand. The antibiotics did not cause the yeast infection. What they did was remove the bacterial community that was preventing it. The *Candida* was already there, held in check by its bacterial competitors. When the competitors were eliminated, the fungus expanded. Understanding this principle — that the microbiome is an ecological system in which bacteria and fungi balance each other — transforms how we think about both antibiotic stewardship and antifungal therapy.

10.7 Gut Fungi and IBD: The Interkingdom Inflammation Loop

A Fungal Fingerprint in Crohn's Disease

The connection between gut fungi and inflammatory bowel disease was first hinted at by an unlikely clue: a blood test. **Anti-*Saccharomyces cerevisiae* antibodies (ASCA)** — IgG and IgA antibodies directed against mannans on the surface of *Saccharomyces cerevisiae* — were identified in the 1990s as a serological marker for Crohn's disease, with sensitivity of approximately 50 to 60 per cent and specificity exceeding 90 per cent. ASCA became a standard component of the clinical workup for IBD, used alongside perinuclear anti-neutrophil cytoplasmic antibodies (pANCA) to distinguish Crohn's disease from ulcerative colitis.

But what did ASCA actually mean? The antibodies were directed against a fungus. Their presence in Crohn's disease implied that the immune system of these patients was mounting an exaggerated response to gut fungi — a response that was somehow linked to the chronic intestinal inflammation that defines the disease. For years, this implication was largely ignored. ASCA was used as a diagnostic tool without much thought given to what it revealed about the role of fungi in IBD pathogenesis.

That changed with the publication of a landmark study by Harry Sokol and colleagues in *Gut* in 2017. Using ITS2 sequencing on faecal samples from 235 IBD patients and 38 healthy controls, Sokol's team produced the first large-scale characterisation of the gut mycobiome in IBD [REF:sokol2017]. Their findings were striking.

The gut mycobiome in IBD was significantly different from that in health. The ratio of **Basidiomycota to Ascomycota** — the two major fungal phyla — was increased in IBD. The proportion of *Saccharomyces cerevisiae* was decreased, while the proportion of *Candida albicans* was increased. Disease-specific alterations in fungal diversity were identified: Crohn's disease, in particular, appeared to create a gut environment that favoured fungi at the expense of bacteria — a finding consistent with the ecological principle that bacterial depletion opens space for fungal expansion.

Most revealingly, Sokol's team performed a concomitant analysis of both bacterial and fungal communities — one of the first studies to do so systematically. In healthy subjects, the correlation network between bacterial and fungal taxa was dense and homogeneous — a web of balanced inter-

kingdom relationships. In IBD patients, this network was dramatically disrupted: the correlations were sparse, unbalanced, and dominated by a few strong pairwise relationships. The healthy interkingdom balance had collapsed.

The Interkingdom Inflammation Loop

These findings, combined with the Iliev Dectin-1 data and the CARD9 studies, suggest a model for how fungi contribute to IBD pathogenesis — a model that centres on the concept of an **interkingdom inflammation loop**.

In a healthy gut, bacteria and fungi coexist in a balanced community. Bacteria suppress fungal overgrowth through competition and metabolite production. Fungi stimulate the immune system through Dectin-1 and CARD9, maintaining anti-fungal readiness without triggering overt inflammation. The immune system tolerates the normal fungal community while remaining prepared to respond if the balance shifts.

In IBD — particularly Crohn's disease — multiple factors converge to disrupt this balance. Genetic susceptibility (including polymorphisms in Dectin-1, CARD9, and other innate immune genes) impairs the immune system's ability to manage commensal fungi. Bacterial dysbiosis — driven by the same genetic, environmental, and dietary factors that promote IBD — weakens the bacterial competitive barrier against fungi. The result is fungal expansion — particularly of *Candida* and, as recent studies have shown, *Malassezia* species — in a gut environment already primed for inflammation.

The expanded fungal community then drives further inflammation through Dectin-1 activation, cytokine production, and adaptive immune responses (including ASCA production). This inflammation worsens the bacterial dysbiosis, weakening the competitive barrier further and permitting additional fungal expansion. The loop self-reinforces: bacterial loss □ fungal expansion □ immune activation □ inflammation □ further bacterial loss.

This model does not claim that fungi *cause* IBD. The genetic, environmental, and immunological factors that initiate the disease are complex and multifactorial. But it does suggest that fungi are active participants in the inflammatory process — amplifiers and perpetuators of a dysbiotic cycle that involves both kingdoms. Ignoring the mycobiome, in this context, means ignoring a significant driver of the disease process.

Where This Matters: If you are a patient with Crohn's disease, there is a reasonable chance that your immune system is producing antibodies against gut fungi — antibodies that your gastroenterologist may have measured as part of your diagnostic workup. Those antibodies are not a laboratory curiosity. They are a sign that the fungal component of your gut ecosystem is abnormal and that your immune system is responding to it. The interkingdom inflammation loop — in which bacterial dysbiosis, fungal expansion, and immune activation feed on each other — is an emerging framework for understanding why IBD is so difficult to treat with single-target therapies. The disease involves not one kingdom of microbes but two, and effective therapy may ultimately need to address both.

10.8 The Skin Mycobiome: *Malassezia* in Seborrheic Dermatitis, Dandruff, and Beyond

A Fungus in Every Flake

We introduced *Malassezia* earlier in this chapter as the dominant fungal genus on human skin. Now we examine its clinical significance more closely — because *Malassezia* is implicated in some of the most common dermatological conditions on Earth.

Dandruff — the white or yellowish flakes of dead skin that shed from the scalp — affects roughly half of the post-pubertal global population. It is so common that it is often dismissed as a cosmetic nuisance rather than a disease. But dandruff is the mild end of a spectrum that extends through **seborrheic dermatitis** (red, scaly, itchy patches on the scalp, face, and chest) to severe inflammatory skin conditions. And the central player in this spectrum is *Malassezia*.

The mechanism is not simply fungal infection. *Malassezia* species are present on virtually everyone's scalp, whether or not they have dandruff. The difference lies in the host's inflammatory response to the fungal metabolites. *Malassezia* species produce **lipases** that break down the triglycerides in sebum into free fatty acids. One of these — **oleic acid** — triggers an inflammatory and hyperproliferative response in the scalp epidermis of susceptible individuals, leading to increased skin cell turnover, impaired barrier function, and the visible flaking that we recognise as dandruff.

The evidence for this mechanism is robust. Antifungal treatments — particularly ketoconazole shampoo — are effective at reducing dandruff, which would not be the case if the condition were purely a disorder of skin cell turnover. Conversely, applying oleic acid to the scalps of dandruff-susceptible individuals can reproduce the flaking response, even in the absence of live fungi. The disease is not caused by the fungus per se, but by the interaction between fungal metabolism and host susceptibility — a theme that recurs throughout mycobiome biology [REF:jung2023].

Beyond the Scalp

Malassezia's clinical significance extends beyond dandruff and seborrheic dermatitis. The genus is implicated in **atopic dermatitis** (eczema), particularly the head-and-neck variant that is common in adults. Patients with atopic dermatitis frequently produce IgE antibodies against *Malassezia* antigens, and the fungal community on their skin is altered compared with healthy individuals — changes that include not just shifts in *Malassezia* species composition but broader mycobiome disruptions [REF:jung2023].

In neonatal intensive care units, *Malassezia furfur* and *Malassezia pachydermatis* have been identified as causes of bloodstream infections in low-birth-weight infants receiving lipid-containing parenteral nutrition. The lipophilic yeasts, normally confined to the skin surface, can colonise intravenous catheters — particularly when the catheter is infusing the very lipid emulsions on which *Malassezia* thrives. This is a stark reminder that a commensal organism, in the wrong context, can become a life-threatening pathogen.

And then there is *Malassezia*'s unexpected appearance in the gut. The finding by Limon and colleagues that *Malassezia restricta* is enriched in the gut mucosa of Crohn's disease patients — and that it exacerbates colitis in an animal model — challenged the assumption that *Malassezia* is exclusively a skin organism. How *Malassezia* reaches the gut (swallowing of skin organisms? direct colonisation?) and why it thrives there in the context of IBD (reduced bacterial competition? altered lipid availability?) remain open questions. But the finding reinforces a lesson we have seen

repeatedly in this book: microbes do not respect the boundaries we draw around body sites. The skin mycobiome and the gut mycobiome are connected — by the immune system, by microbial migration, and by the ecological principles that govern both.

10.9 The Infant Mycobiome: Colonisation, Maturation, and Early Immune Priming

The Forgotten Founders

In Chapter 3, we described the dramatic process by which the newborn gut is colonised by bacteria — the microbial baptism of birth, the selective feeding of *Bifidobacterium* by human milk oligosaccharides, the gradual maturation from a simple pioneer community to the complex ecosystem of the adult gut. But that account, like most accounts of early-life microbiome development, was entirely about bacteria. What about fungi?

The infant gut mycobiome remains far less well characterised than its bacterial counterpart, but the picture that is emerging is both surprising and clinically significant.

Fungal colonisation begins at birth — or possibly before, though the evidence for prenatal fungal exposure is even more limited and contested than the evidence for prenatal bacterial exposure. Vaginal delivery exposes the newborn to the maternal vaginal mycobiome, which is typically dominated by *Candida* species. Caesarean delivery, by bypassing the vaginal canal, alters the initial fungal inoculum — a difference that, as a 2022 study by Wang and colleagues demonstrated, has measurable effects on the skin mycobiome of children that persist for years. Children born by caesarean section exhibited a mycobiome assembled by more deterministic, niche-based processes and displayed a more fragile microbial network compared with vaginally born children [REF:wang2022].

In the gut, the earliest fungal colonisers are typically *Candida* species — particularly *C. albicans* and *C. parapsilosis* — which are acquired from the mother and the hospital environment. Over the first months of life, the fungal community diversifies, with the appearance of *Saccharomyces*, *Malassezia*, *Rhodotorula*, and other genera. The trajectory of this diversification is influenced by many of the same factors that shape bacterial colonisation: delivery mode, feeding method (breast milk vs formula), antibiotic exposure, and the microbial environment of the home.

Immune Consequences of Early Fungal Colonisation

A 2017 review by Tonya Ward, Dan Knights, and Cheryl Gale in *BMC Medicine* drew attention to a critical gap in microbiome research: the almost complete absence of studies examining the immunological consequences of early fungal colonisation [REF:ward2017]. The bacterial microbiome's role in immune development — in training the infant immune system to distinguish self from non-self, in shaping the balance between pro-inflammatory and regulatory immune responses — had been extensively studied. The fungal microbiome's equivalent role was essentially unknown.

Since that review, several studies have begun to fill the gap. A 2022 review by Gutierrez, van Tilburg Bernardes, and Arrieta synthesised the emerging evidence and proposed a framework for understanding how the intestinal mycobiome shapes mucosal and systemic immunity from birth [REF:gutierrez2022]. The key insight is that early fungal colonisation — like early bacterial colonisation — provides critical immune training. The infant immune system must learn to tolerate com-

mensal fungi while maintaining the capacity to respond to fungal invasion. The Dectin-1/CARD9 pathway, discussed in section 10.5, is central to this learning process.

Disruptions to early fungal colonisation — whether through caesarean delivery, antibiotic exposure, or antifungal treatment — may therefore have immunological consequences that parallel the consequences of disrupted bacterial colonisation. There is preliminary evidence linking early-life mycobiome alterations to the development of allergic diseases, though the data remain sparse and the causal mechanisms uncertain.

A striking 2025 study by Gutierrez, Arrieta, and colleagues in *Nature Communications* provided some of the first causal evidence [REF:gutierrez2025]. Using gnotobiotic mice colonised with three core species of the infant gut mycobiome — *Rhodotorula mucilaginosa*, *Malassezia restricta*, and *Candida albicans* — they showed that each species had distinct effects on metabolic health. *R. mucilaginosa* and *M. restricta* increased adiposity, while *C. albicans* conferred leanness and resistance to diet-induced obesity. Each species also distinctly impacted the immune landscape in white adipose tissue, suggesting that early fungal colonisation shapes metabolic health through immune modulation. The work revealed that common fungal colonisers are not interchangeable — their specific identities matter for long-term host physiology.

The Missing Piece

The infant mycobiome represents one of the most significant gaps in current microbiome research. We know that fungi colonise the infant gut from birth. We know that delivery mode and feeding method influence the trajectory of fungal colonisation. We know that the immune system has dedicated pathways for recognising and responding to fungi. And we have tantalising evidence that early fungal colonisation influences immune development and metabolic health.

What we lack is the kind of large-scale, longitudinal, multi-kingdom study that has transformed our understanding of bacterial colonisation. The bacterial story — from the Human Microbiome Project to the TEDDY study — has been told with increasing clarity and detail over the past two decades. The fungal story is still in its opening chapter. Filling this gap is not merely an academic exercise. If early fungal colonisation proves to be as immunologically important as early bacterial colonisation — and the evidence increasingly suggests that it is — then the interventions we develop to support healthy microbiome development will need to consider not just bacteria, but the forgotten kingdom as well.

In the next chapter, we venture into the most controversial territory in microbiome science: the question of whether microbes reside in the brain itself. The central nervous system — long held to be sterile, protected by the blood-brain barrier from microbial intrusion — may harbour its own microbial community. Chapter 11 examines the evidence, the scepticism, and the extraordinary implications of the proposed brain microbiome.

Chapter 10 — Key References

- [REF:iliev2012] Iliev, I.D., Funari, V.A., Taylor, K.D., Nguyen, Q., Reyes, C.N., Strom, S.P., Brown, J., Becker, C.A., Fleshner, P.R., Dubinsky, M., Rotter, J.I., Wang, H.L., McGovern, D.P.B., Brown, G.D. & Underhill, D.M. (2012). Interactions between commensal fungi and

the C-type lectin receptor Dectin-1 influence colitis. *Science*, 336(6086), 1314–1317. DOI: 10.1126/science.1221789

- [REF:zhang2022] Zhang, F., Aschenbrenner, D., Yoo, J.Y. & Zuo, T. (2022). The gut mycobiome in health, disease, and clinical applications in association with the gut bacterial microbiome assembly. *The Lancet Microbe*, 3(12), e969–e983. DOI: 10.1016/S2666-5247(22)00203-8
- [REF:findley2013] Findley, K., Oh, J., Yang, J., Conlan, S., Deming, C., Meyer, J.A., Schoenfeld, D., Nomicos, E., Park, M., Kong, H.H. & Segre, J.A. (2013). Topographic diversity of fungal and bacterial communities in human skin. *Nature*, 498(7454), 367–370. DOI: 10.1038/nature12171
- [REF:iliev2013] Iliev, I.D. & Underhill, D.M. (2013). Striking a balance: fungal commensalism versus pathogenesis. *Current Opinion in Microbiology*, 16(3), 366–373. DOI: 10.1016/j.mib.2013.05.004
- [REF:doron2021] Doron, I., Leonardi, I., Li, X.V., Fiers, W.D., Semon, A., Bialt-DeCelie, M., Migaud, M., Gao, I.H., Lin, W.-Y., Kusakabe, T., Puel, A. & Iliev, I.D. (2021). Human gut mycobiota tune immunity via CARD9-dependent induction of anti-fungal IgG antibodies. *Cell*, 184(4), 1017–1031.e14. DOI: 10.1016/j.cell.2021.01.016
- [REF:sokol2017] Sokol, H., Leducq, V., Aschard, H., Pham, H.-P., Jegou, S., Landman, C., Cohen, D., Liguori, G., Bourrier, A., Nion-Larmurier, I., Cosnes, J., Seksik, P., Langella, P., Skurnik, D., Richard, M.L. & Beaugerie, L. (2017). Fungal microbiota dysbiosis in IBD. *Gut*, 66(6), 1039–1048. DOI: 10.1136/gutjnl-2015-310746
- [REF:noverr2004] Noverr, M.C., Noggle, R.M., Toews, G.B. & Huffnagle, G.B. (2004). Role of antibiotics and fungal microbiota in driving pulmonary allergic responses. *Infection and Immunity*, 72(9), 4996–5003. DOI: 10.1128/IAI.72.9.4996-5003.2004
- [REF:noverr2005] Noverr, M.C., Falkowski, N.R., McDonald, R.A., McKenzie, A.N. & Huffnagle, G.B. (2005). Development of allergic airway disease in mice following antibiotic therapy and fungal microbiota increase: role of host genetics, antigen, and interleukin-13. *Infection and Immunity*, 73(1), 30–38. DOI: 10.1128/IAI.73.1.30-38.2005
- [REF:jung2023] Jung, W.H. (2023). Alteration in skin mycobiome due to atopic dermatitis and seborrheic dermatitis. *Biophysics Reviews*, 4(1), 011309. DOI: 10.1063/5.0136543
- [REF:belvoncikova2022] Belvončíková, P., Splichalova, P., Videnska, P. & Gardlik, R. (2022). The human mycobiome: colonization, composition and the role in health and disease. *Journal of Fungi*, 8(10), 1046. DOI: 10.3390/jof8101046
- [REF:wang2022] Wang, Y.-R., Zhu, T., Kong, F.-Q., Duan, Y.-Y., Galzote, C. & Quan, Z.-X. (2022). Infant mode of delivery shapes the skin mycobiome of prepubescent children. *Microbiology Spectrum*, 10(5), e02267-22. DOI: 10.1128/spectrum.02267-22
- [REF:ward2017] Ward, T.L., Knights, D. & Gale, C.A. (2017). Infant fungal communities: current knowledge and research opportunities. *BMC Medicine*, 15(1), 30. DOI: 10.1186/s12916-017-0802-z
- [REF:gutierrez2022] Gutierrez, M.W., van Tilburg Bernardes, E., Changirwa, D., McDonald, B. & Arrieta, M.-C. (2022). “Molding” immunity — modulation of mucosal and systemic im-

munity by the intestinal mycobiome in health and disease. *Mucosal Immunology*, 15(4), 573–583. DOI: 10.1038/s41385-022-00515-w

- [REF:gutierrez2025] Gutierrez, M.W., van Tilburg Bernardes, E., Ren, E., Kalbfleisch, K.N., Day, M., Lameu, E.L., Glatthardt, T., Mercer, E.M., Sharma, S., Zhang, H., Al-Azawy, A., Chleilat, F., Hirota, S.A., Reimer, R.A. & Arrieta, M.-C. (2025). Early-life gut mycobiome core species modulate metabolic health in mice. *Nature Communications*, 16(1), 1467. DOI: 10.1038/s41467-025-56743-8
- [REF:he2025] He, S., Sun, Y., Yu, J., Tang, M., Zhang, Q., Meng, B., Fan, R., Liu, Z., Liu, Y., Hu, L., Wu, T. & Li, J. (2025). Dysbiotic gut fungi exacerbate *Klebsiella pneumoniae* lung infection via Dectin-1-mediated alveolar macrophage hyperactivation. *The ISME Journal*, 19(1). DOI: 10.1093/ismejo/wraf181
- [REF:hsu2023] Hsu, C., Ghannoum, M., Cominelli, F. & Di Martino, L. (2023). Mycobiome and inflammatory bowel disease: role in disease pathogenesis, current approaches and novel nutritional-based therapies. *Inflammatory Bowel Diseases*, 29(3), 470–479. DOI: 10.1093/ibd/izac156

Additional references to be verified and expanded during manuscript preparation.

Chapter 11: The Brain Microbiome — The Final Frontier

In every chapter of this book so far, we have followed the same narrative arc: a body site long assumed to be sterile turns out to harbour microbial life. The lungs, the urinary tract, the uterus, the lower respiratory tract — one by one, the “sterile” compartments of the body have fallen. In each case, the pattern has been remarkably consistent. Initial reports of microbial detection are met with scepticism. The sceptics invoke contamination. Better-controlled studies follow. Eventually, a consensus emerges: the microbes are real, they are there, and they matter.

Now we arrive at the last and most heavily fortified bastion of sterility: the brain.

The central nervous system occupies a unique position in the body’s defensive architecture. It is shielded by the **blood-brain barrier (BBB)** — a selectively permeable boundary formed by specialised endothelial cells, joined by exceptionally tight junctions, and reinforced by astrocyte foot processes and pericytes — that excludes the vast majority of molecules and virtually all microorganisms from the brain parenchyma. The BBB is not merely a passive wall. It is an active, regulated interface that permits the entry of glucose, amino acids, and certain lipophilic molecules while rejecting pathogens, toxins, and most drugs. It is, by design, one of the most effective biological barriers in nature.

For over a century, the logical consequence of this barrier was taken as axiomatic: the healthy brain is sterile. Microbes in the brain meant infection — meningitis, encephalitis, brain abscess — and infection meant disease. The idea that bacteria, fungi, or viruses might reside in healthy brain tissue as commensals, the way they reside in the gut or on the skin, was not merely unproven. It was, for most neuroscientists and microbiologists, unthinkable.

And yet, over the past decade, a series of studies has challenged this assumption. Bacteria have been detected in postmortem brain tissue — not only from patients with neurodegenerative disease but from neurologically normal controls. Fungi have been visualised in brain sections from Alzheimer’s patients. Viral DNA has been found woven into the molecular networks of the ageing brain. And in 2018, a presentation at the Society for Neuroscience meeting in San Diego produced electron microscopy images that appeared to show bacteria sitting comfortably inside the star-shaped astrocytes of the human brain.

This chapter examines the evidence for and against a brain microbiome — the most controversial proposition in contemporary microbiome science. We will proceed carefully, because the territory demands it. The evidence is genuinely ambiguous. The methodological pitfalls are severe. The implications, if the hypothesis is even partially correct, are enormous. And the history of microbiome research — in which “sterile” has repeatedly turned out to mean “not yet looked at properly” — should make us cautious about dismissing the possibility out of hand.

11.1 The Dogma of the Sterile Brain and How It Fell

The Fortress

The idea that the brain is immunologically privileged — sealed off from the microbial world that teems through the rest of the body — has deep roots in medical science. The concept dates to experiments performed in the 1920s by the Japanese scientist Y. Shirai, who showed that tumour tissue transplanted into the brain of rats was not rejected by the immune system, whereas the same tissue transplanted elsewhere was rapidly destroyed. The interpretation was that the brain existed behind an immunological curtain: a space where the normal rules of immune surveillance did not apply.

The blood-brain barrier itself had been recognised even earlier. In the late nineteenth century, Paul Ehrlich — the same Paul Ehrlich who coined the term “chemotherapy” and won the Nobel Prize for his work on immunity — observed that certain dyes injected into the bloodstream stained every organ in the body except the brain. His student Edwin Goldmann later showed the reverse: dye injected directly into the cerebrospinal fluid stained the brain but not the rest of the body. Something was keeping the brain’s fluid compartments separate from the general circulation. That something, eventually characterised at the ultrastructural level in the 1960s by Thomas Reese and Morris Karnovsky using electron microscopy, was the tight-junction network of the cerebral capillary endothelium — the physical basis of the BBB [REF:reese1967].

The BBB is not a single structure but a multi-layered system. The innermost layer — the endothelial cells of the brain’s capillaries — is joined by **tight junctions** (also called *zonulae occludentes*) that are far more restrictive than the junctions between endothelial cells elsewhere in the body. In peripheral capillaries, small molecules and even some proteins can pass between endothelial cells through gaps and fenestrations. In brain capillaries, these gaps are sealed. The endothelial cells are further enveloped by **pericytes** — contractile cells embedded in the capillary basement membrane — and by the **end-feet of astrocytes**, star-shaped glial cells that wrap around the capillary surface and regulate ion and water transport. Together, these three cell types form a tripartite barrier that is exquisitely selective: it admits glucose (via GLUT1 transporters), certain amino acids (via specific carrier systems), and a small number of lipophilic molecules that can dissolve through the cell membranes, while excluding bacteria, viruses, fungi, most proteins, and the vast majority of drugs.

The clinical significance of this barrier is immense. It is the reason that brain infections, when they do occur, are medical emergencies: once a pathogen breaches the BBB, the brain’s own immune resources — primarily microglia, the resident macrophages of the central nervous system — must mount a defence in a confined space where inflammation itself can cause devastating damage. It is also the reason that treating brain infections is so difficult: most antibiotics cannot cross the BBB, and those that can — such as the lipophilic fluoroquinolones and certain β -lactams administered at high doses — often achieve only marginal concentrations in the cerebrospinal fluid.

For generations of physicians and scientists, the strength of this barrier was taken as proof that the healthy brain was microbe-free. The logic was straightforward: if the BBB keeps microbes out, and the brain shows no signs of infection, then the brain must be sterile. It was sound reasoning — as far as it went. But as we have seen repeatedly in this book, the absence of detected microbes is not the same as the absence of microbes. It depends entirely on whether anyone has looked, with what tools, and with what expectations.

The First Cracks

The dogma of the sterile brain was never quite as absolute as it appeared. Individual pathogens had been found in brain tissue for over a century — *Treponema pallidum* in neurosyphilis, *Mycobacterium tuberculosis* in tuberculous meningitis, *Toxoplasma gondii* in immunocompromised patients, various herpesviruses in encephalitis. These were understood as infections: breaches of the barrier that caused recognisable disease. They did not challenge the sterility dogma because they were pathological, not commensal. The microbes were invaders, not residents.

The first challenge to the paradigm came not from bacteria but from viruses — and not from acute infection but from latency. **Herpes simplex virus type 1 (HSV-1)**, the virus responsible for cold sores, is a neurotropic virus that, after primary infection, establishes lifelong latency in the trigeminal ganglia — nerve cell bodies that sit just outside the brain. In the latent state, the virus produces no infectious particles and causes no symptoms. But it persists, silently, in neurons — and periodic reactivation can send viral particles along nerve fibres toward the brain. The question of whether HSV-1 might also establish latency within the brain itself, rather than merely in the peripheral ganglia, was raised as early as the 1980s by Ruth Itzhaki and colleagues at the University of Manchester, who detected HSV-1 DNA in postmortem brain tissue from both Alzheimer's patients and elderly controls [REF:jamieson1991].

Itzhaki's finding was important for two reasons. First, it demonstrated that viral DNA could be present in brain tissue without causing overt encephalitis — that a virus could be *in* the brain without the brain being *infected* in the clinical sense. Second, it raised the possibility that this latent viral presence might not be entirely benign. Itzhaki proposed that HSV-1, in combination with the APOE-ε4 allele (the strongest genetic risk factor for late-onset Alzheimer's disease), might contribute to the neuroinflammation and amyloid accumulation that characterise Alzheimer's pathology. This **viral hypothesis of Alzheimer's disease** would remain controversial for three decades — but the fundamental observation, that viral genetic material could reside in apparently normal brain tissue, planted a seed of doubt about the sterility assumption.

The bacterial challenge came later. In 1998, Brian Balin and colleagues at the Philadelphia College of Osteopathic Medicine reported detecting ***Chlamydia pneumoniae*** — an intracellular bacterium best known as a cause of respiratory infections — in postmortem brain tissue from patients with late-onset Alzheimer's disease [REF:balin1998]. Using PCR, immunohistochemistry, and electron microscopy, the team found *C. pneumoniae* in 17 of 19 Alzheimer's brains, localised to areas of neuropathology — particularly near amyloid plaques and neurofibrillary tangles. In contrast, only 1 of 19 age-matched control brains was positive. The bacterium appeared to be present within microglia, astrocytes, and perivascular macrophages — the very immune cells that patrol the brain's internal environment.

The response from the scientific community was cautious. A replication attempt by Gieffers and colleagues in Germany, published in 2000, failed to detect *C. pneumoniae* DNA in brain tissue from 20 Alzheimer's patients using nested PCR and immunohistochemistry [REF:gieffers2000]. Other groups reported mixed results. The inconsistency was troubling, and it raised a question that would haunt every subsequent study of microbes in the brain: were the positive results real, or were they artefacts of contamination?

11.2 Evidence for Microbial Presence in the CNS — What We Know So Far

Bacteria in the Brain

The question of whether bacteria inhabit the human brain moved from the margins to the mainstream in November 2018, when Rosalinda Roberts and colleagues from the University of Alabama at Birmingham (UAB) presented a poster at the annual meeting of the Society for Neuroscience in San Diego. The images were striking: high-resolution electron microscopy photographs showing rod-shaped bacteria residing inside astrocytes and in close proximity to the blood-brain barrier in postmortem human brain tissue. The bacteria appeared intact, not degraded, and were present in samples from both schizophrenia patients and neurologically normal controls.

Roberts' team had not set out to find bacteria. They had been studying ultrastructural differences in the brains of schizophrenia patients when they noticed the bacterial profiles — consistent, reproducible, and present across multiple brain regions. To address the obvious contamination concern, they performed a critical control experiment: they prepared brain tissue from germ-free mice — animals raised in completely sterile conditions with no microbial exposure — using the same fixation and processing protocols. The germ-free mouse brains showed no bacteria. Conventionally raised mouse brains, processed identically, did.

The presentation generated immediate attention and immediate scepticism. The bacteria had been detected by electron microscopy, not by sequencing or culture, which meant that their identity was unknown. The postmortem interval — the time between death and tissue preservation — was a concern, because bacteria from the gut or bloodstream might translocate to the brain after death. And the history of microbiome research was littered with examples of contamination artefacts that had initially appeared convincing. The germ-free mouse control was reassuring, but not definitive: it demonstrated that the processing protocol did not introduce bacteria, but it did not rule out the possibility that bacteria had entered the brain postmortem before fixation.

Roberts' work remained unpublished as a full peer-reviewed paper for several years, but it catalysed a new wave of investigation. Meanwhile, other groups had been approaching the same question from different angles.

The Sequencing Evidence

The application of 16S rRNA gene sequencing to brain tissue had begun before the Roberts presentation. In 2017, David Emery, Shelley Allen, and colleagues at the University of Bristol published the first study to apply next-generation 16S sequencing to postmortem human brain tissue [REF:emery2017]. Comparing Alzheimer's patients with cognitively unimpaired controls, they found bacterial DNA in brain samples from both groups — but with a significantly higher abundance in the Alzheimer's brains. The dominant phyla were the same ones that dominate the gut and other body sites: Proteobacteria, Firmicutes, Actinobacteria, and Bacteroidetes.

In 2020, Susan Westfall, Duy Dinh, and Giulio Maria Pasinetti at the Icahn School of Medicine at Mount Sinai extended this work in a study that explicitly addressed the problem of study bias [REF:westfall2020]. They collected postmortem brain samples from two independent cohorts processed at different locations with different extraction protocols, and applied V4 16S sequencing to hippocampal and cerebellar tissue. Their findings were nuanced. Bacterial DNA was detectable in brain samples from both AD and control subjects. The dominant phyla — Proteobacteria, Firmicutes, Actinobacteria, Bacteroidetes — were consistent across cohorts. There were some differences between brain regions (hippocampus versus cerebellum) and between AD and control groups, but

the most striking finding was how much the results varied *between* the two cohorts. Alpha and beta diversity differed significantly depending on which cohort the samples came from — suggesting that technical variables in tissue collection and processing had a powerful influence on the detected microbiome. The postmortem interval, in particular, appeared to affect bacterial DNA recovery, raising the spectre that at least some of the detected bacteria represented postmortem invasion rather than antemortem residence.

The Westfall study was admirably cautious in its conclusions: “Regardless of the study bias, this study concludes that bacterial DNA can be isolated from the human brain suggesting that a brain microbiome may exist; however, more studies are required to understand the variation.”

A more ambitious attempt came in 2023, when Yves Moné, Garth Ehrlich, and colleagues at Drexel University College of Medicine published a study using full-length 16S rRNA gene sequencing with PacBio long-read technology — a method offering far higher taxonomic resolution than the short-read approaches used in earlier studies [REF:mone2023]. They analysed 130 samples from 32 subjects (16 AD, 16 control), spanning the frontal lobe, temporal lobe, and entorhinal cortex. They detected bacteria in both cohorts, with the principal taxa including *Cutibacterium acnes* (formerly *Propionibacterium acnes*) — the common skin commensal — and species of *Acinetobacter* and *Comamonas*. Using a sophisticated computational approach borrowed from computational linguistics (the Latent Dirichlet Allocation algorithm), they identified five distinct microbiota classes across the samples and proposed a temporal model in which the brain microbiome composition changed over time, with pathogenicity associated not with any single species but with a complex polymicrobial dynamic.

The Moné study was notable for its methodological rigour and its intellectual ambition, but it also highlighted a recurring problem: the bacteria most frequently detected in brain tissue — *Cutibacterium acnes*, *Acinetobacter*, various Proteobacteria — are also among the most common contaminants in molecular biology reagents and laboratory environments. This overlap between “brain bacteria” and “reagent contaminants” is the central methodological challenge that has made the brain microbiome question so difficult to resolve.

Fungi in the Brain

The evidence for fungi in brain tissue comes primarily from the laboratory of Luis Carrasco at the Autonomous University of Madrid, who published a series of studies between 2014 and 2018 reporting fungal detection in postmortem brain samples from patients with Alzheimer’s disease, amyotrophic lateral sclerosis, and other neurodegenerative conditions [REF:alonso2014, REF:pisa2015].

Using a combination of immunohistochemistry with anti-fungal antibodies, PCR with fungal-specific primers, and proteomic analysis, Carrasco’s group reported detecting fungal cells and macromolecules in brain tissue from Alzheimer’s patients across multiple brain regions — including the frontal cortex, hippocampus, and entorhinal cortex. The fungal genera most frequently identified were *Candida*, *Malassezia*, *Saccharomyces*, and *Alternaria* — organisms familiar to us from previous chapters as common residents of the gut and skin. Control brains showed either no fungal detection or significantly lower levels.

In 2015, the group published a particularly striking paper in *Scientific Reports* titled “Different Brain Regions are Infected with Fungi in Alzheimer’s Disease,” in which they used confocal microscopy to visualise fungal structures — including intracellular yeast cells and extracellular

hyphal-like filaments — in brain sections stained with anti-chitin and anti-fungal antibodies [REF:pisa2015]. The images were arresting: green-fluorescent structures, clearly distinct from the surrounding neural tissue, distributed across the parenchyma and within blood vessel walls.

These findings have been received with a mixture of interest and caution. The Carrasco group's work has been criticised for relying heavily on immunohistochemistry — a technique susceptible to non-specific antibody binding and cross-reactivity — and for the absence of germ-free or adequately processed negative controls in some studies. Furthermore, the finding of *Malassezia* (a skin commensal) and *Alternaria* (an environmental mould) in brain tissue raises the same contamination concerns that plague the bacterial studies. On the other hand, the consistency of the findings across multiple patients, multiple brain regions, and multiple detection methods gives them a persistence that is difficult to dismiss entirely.

Viruses: The Strongest Case

If the evidence for a bacterial brain microbiome is contested and the evidence for a fungal brain presence is preliminary, the case for viral genetic material in the brain is considerably stronger — in large part because the biology of viral latency provides a plausible mechanism for how viruses could get there and stay there without causing acute disease.

We have already mentioned Itzhaki's work on HSV-1. In 2018, a landmark computational study by Ben Readhead and colleagues, published in *Neuron*, brought the viral-brain connection to a new level of sophistication [REF:readhead2018]. Readhead's team performed a multiscale analysis of transcriptomic, genomic, proteomic, and neuropathological data from three independent Alzheimer's disease cohorts — the Mount Sinai Brain Bank, the Mayo Clinic, and the Religious Orders Study/Memory and Aging Project. They identified increased abundance of human herpesvirus 6A (HHV-6A) and human herpesvirus 7 (HHV-7) in AD brains compared with controls, and found that viral abundance was correlated with the expression of genes involved in amyloid precursor protein (APP) metabolism — including *BACE1*, *PSEN1*, *CLU*, and *BIN1*, all established Alzheimer's risk genes.

The implications were provocative: not merely that viral DNA was present in the AD brain, but that the viruses appeared to be integrated into the molecular networks driving AD pathology. The viruses were not passive passengers. They appeared to be actively modulating gene expression in ways that could promote amyloid production and neurodegeneration.

However, the Readhead findings were subsequently challenged from multiple directions. A computational reanalysis by Samuel Chorlton, published in 2020, re-examined the same sequencing datasets using alternative bioinformatic pipelines and found no HHV-7 reads and only minimal HHV-6A reads in the samples with the highest reported viral abundance [REF:chorlton2020]. Independently, Mary Alice Allnutt and colleagues screened for HHV-6 across three independent AD brain repositories using both RNA-seq data and DNA-based PCR, finding no statistically significant difference in HHV-6A or HHV-6B abundance between AD and control brains [REF:allnutt2020]. The discrepancies appeared to stem from differences in how viral sequences were identified and quantified — highlighting the sensitivity of metagenomics results to computational and methodological choices.

Meanwhile, a parallel line of investigation was producing some of the most intriguing findings in the field. In 2018, William Eimer, Robert Moir, Rudolph Tanzi, and colleagues at Massachusetts General Hospital published a study in *Neuron* demonstrating that **amyloid-beta (A β) peptides**

— the protein fragments that aggregate to form the hallmark plaques of Alzheimer’s disease — function as antimicrobial peptides [REF:eimer2018]. They showed that A β oligomers bind herpesvirus surface glycoproteins, accelerating β -amyloid deposition and leading to protective viral entrapment in both mouse and human neural cell culture models. In effect, the brain appeared to be deploying amyloid as a defence against viral infection — and the plaques that define Alzheimer’s disease might be the collateral damage of an antimicrobial response gone awry.

This **antimicrobial protection hypothesis** — the idea that A β is not a meaningless waste product but an innate immune effector that seeds amyloid in response to microbial challenge — has become one of the most actively investigated frameworks in Alzheimer’s research. If correct, it would transform our understanding of the disease: Alzheimer’s would not be primarily a disorder of protein misfolding but an inflammatory response to microbial presence in the brain, with amyloid plaques as the immunological scar tissue.

11.3 Methodological Controversies: Contamination Artefact or Genuine Residents?

The Contamination Problem

Every study of the brain microbiome faces a fundamental challenge: how do you distinguish a genuine resident from a contaminant?

The problem is not unique to brain research. As we discussed in Chapter 4, the “**kitome**” — the microbial DNA present in molecular biology reagents, DNA extraction kits, and laboratory consumables — is a pervasive source of false-positive signals in low-biomass microbiome studies. The term was coined by Susannah Salter and colleagues at the Wellcome Sanger Institute, who published a landmark paper in 2014 demonstrating that DNA extraction kits contain a characteristic community of bacterial DNA — dominated by Proteobacteria, Actinobacteria, and Firmicutes — that can overwhelm the genuine microbial signal in samples with low microbial content [REF:salter2014]. The kitome taxa identified by Salter’s team overlap substantially with the bacteria reported in brain tissue studies: *Bradyrhizobium*, *Ralstonia*, *Propionibacterium* (now *Cutibacterium*), *Acinetobacter*, and various unclassified Proteobacteria.

This overlap is the crux of the contamination argument. Critics of the brain microbiome hypothesis point out that:

The brain, even if it does harbour microbes, would be an extremely low-biomass environment — perhaps containing only a few hundred or thousand bacterial cells per gram of tissue, compared with roughly 10^{11} cells per gram in the colon. At such low biomass levels, even trace amounts of contaminating DNA can dominate the sequencing results.

The most commonly reported “brain bacteria” are the same taxa that populate the kitome. This does not prove they are contaminants, but it makes it impossible to distinguish real signal from noise without extraordinarily rigorous controls.

Many brain microbiome studies have used postmortem tissue with varying postmortem intervals. After death, the blood-brain barrier breaks down, and gut bacteria can translocate throughout the body — a process well documented in forensic microbiology. A study that finds gut-associated bacteria in brain tissue collected hours after death cannot easily distinguish antemortem colonisation from postmortem invasion.

Christopher Link, in his 2021 review paper “Is There a Brain Microbiome?” published in *Neuroscience Insights*, crystallised the sceptical position with admirable clarity [REF:link2021]. He argued that “the evidence for the presence of microbes in diseased brains is quite strong, but a compelling demonstration of resident microbes in the healthy human brain remains to be done.” To his knowledge, no study had yet looked specifically for evidence of a healthy brain microbiome using unbiased and appropriately controlled methods. The key word was *appropriately* — because adequate controls for a brain microbiome study would need to include: (1) negative extraction controls processed alongside every brain sample; (2) no-template PCR controls; (3) statistical comparison of sample communities against control communities using established decontamination algorithms; and (4) ideally, tissue from germ-free animals processed with identical protocols.

The Counter-Arguments

Proponents of the brain microbiome hypothesis offer several responses to the contamination critique.

First, the Roberts electron microscopy data are not susceptible to kitome contamination, because electron microscopy detects physical structures, not DNA. The bacteria observed in brain tissue by Roberts’ team were morphologically intact, located within specific cell types (astrocytes), and absent from germ-free mouse controls processed identically. While electron microscopy cannot identify the species of bacterium, it can confirm that something bacterial — with the characteristic double-membrane structure, size, and morphology of a gram-negative rod — is physically present in the tissue. DNA contamination cannot produce an electron-dense structure inside an astrocyte.

Second, several studies have reported that the composition of microbial communities detected in brain tissue varies systematically by brain region — with, for example, the hippocampus harbouring a different community than the cerebellum [REF:westfall2020]. If the detected microbes were purely contaminants, they would be expected to be distributed uniformly across all brain regions (since all regions are processed with the same kits). Regional variation, while not definitive proof of genuine colonisation, is at least consistent with it.

Third, the detection of microbial components by multiple independent methods — 16S sequencing, immunohistochemistry, electron microscopy, proteomics — provides a degree of cross-validation. Each method has its own vulnerabilities, but the specific failure modes are different. Kitome contamination affects sequencing but not electron microscopy. Non-specific antibody binding affects immunohistochemistry but not 16S sequencing. When multiple methods converge on the same conclusion, the probability that all of them are simultaneously producing artefacts decreases.

Fourth, and perhaps most importantly, the immune system of the brain behaves as though it expects to encounter microbes. Microglia — the brain’s resident immune cells — express toll-like receptors (TLRs), pattern recognition receptors that detect microbial molecules such as lipopolysaccharide (TLR4), bacterial lipoproteins (TLR2), and viral double-stranded RNA (TLR3). If the brain were truly sterile — if microbes never entered the CNS — the evolutionary maintenance of these receptors would be difficult to explain. The simplest explanation is that the brain’s immune system has evolved to deal with a level of microbial exposure that, while far lower than the gut or skin, is not zero.

The State of Play

The honest assessment, as of 2026, is that the existence of a brain microbiome — in the sense of a stable, resident microbial community analogous to the gut or skin microbiome — remains unproven. The evidence is suggestive but not conclusive. What is clear is that microbial molecules — DNA, proteins, cell wall components — can be detected in brain tissue, and that their presence is not always attributable to contamination. Whether these molecules represent living microbes, dead microbes, microbial debris translocated from the gut or bloodstream, or some combination of all three is the central unresolved question.

The field is analogous to where lung microbiome research stood in approximately 2010: tantalising data, legitimate contamination concerns, heated debate, and a growing recognition that the truth probably lies somewhere between “the brain is sterile” and “the brain has a rich microbiome.” The resolution will require studies conducted with the kind of meticulous contamination controls that have become standard in low-biomass microbiome research — studies that, as Link noted, have not yet been adequately performed for the healthy human brain.

11.4 Potential Implications for Neurodegeneration

The Infection Hypothesis of Alzheimer’s Disease

Even before the brain microbiome debate began, a parallel and overlapping line of investigation had been building the case that microbial infection might play a role in the most common neurodegenerative diseases. The **infection hypothesis of Alzheimer’s disease** — the proposal that one or more pathogens contribute to the initiation or progression of AD — has been advanced in various forms for over three decades, and it has accumulated a body of evidence that, while falling short of proof, is too substantial to ignore.

The cast of suspected pathogens is diverse. We have already encountered several:

Herpes simplex virus type 1 (HSV-1): Detected in brain tissue from both AD patients and controls, but associated with increased Alzheimer’s risk in carriers of the APOE-ε4 allele. The Eimer et al. (2018) finding that Aβ functions as an antimicrobial peptide that entraps herpesviruses provides a mechanistic link between viral presence and amyloid pathology [REF:eimer2018].

Human herpesvirus 6A (HHV-6A) and HHV-7: Identified at elevated levels in AD brains by Readhead et al. (2018), with evidence of regulatory relationships between viral abundance and AD-related gene expression — though subsequently disputed by reanalysis [REF:readhead2018, REF:chorlton2020, REF:allnutt2020].

Chlamydia pneumoniae: Detected in AD brain tissue by Balin et al. (1998), with localisation to areas of neuropathology — but with inconsistent replication by other groups [REF:balin1998, REF:gieffers2000].

Porphyromonas gingivalis: The keystone pathogen of chronic periodontitis — gum disease — was the subject of a landmark 2019 paper by Stephen Dominy, Jan Potempa, and colleagues, published in *Science Advances* [REF:dominy2019]. The study identified *P. gingivalis* in the brains of Alzheimer’s patients, detected its toxic proteases (gingipains) in brain tissue at levels that correlated with tau and ubiquitin pathology, demonstrated that oral infection with *P. gingivalis* in mice led to

brain colonisation and increased A β ₁₋₄₂ production, and showed that small-molecule gingipain inhibitors could reduce bacterial brain load and rescue hippocampal neurons in infected mice.

The Dominy study was significant not only for its findings but for the directness of its mechanistic hypothesis: a bacterium from the mouth reaches the brain, produces neurotoxic enzymes, and drives the molecular changes associated with Alzheimer's pathology. The company behind the research, Cortexyme, subsequently advanced a gingipain inhibitor (atuzaginstat, COR388) into clinical trials for Alzheimer's disease — representing one of the first direct tests of the infection hypothesis in human patients. The Phase 2/3 trial (the GAIN trial) ultimately did not meet its primary endpoints, and Cortexyme faced scientific and business setbacks. But the trial itself was a watershed: it marked the first time the pharmaceutical industry had taken the microbial hypothesis of Alzheimer's seriously enough to bet a drug programme on it.

Fungi: The Carrasco group's detection of fungal cells and macromolecules in AD brain tissue adds another kingdom to the list of potential microbial contributors [REF:alonso2014, REF:pisa2015]. If confirmed, the presence of fungi in the brain would have particular significance given what we learned in Chapter 10: fungal cell wall components, particularly β -glucan, are among the most potent activators of innate immunity, and the Dectin-1/CARD9 pathway that detects them triggers robust inflammatory responses. Chronic, low-level fungal stimulation of brain-resident immune cells could, in principle, contribute to the sustained neuroinflammation that characterises AD.

The Common Thread: Neuroinflammation

What connects these diverse pathogens — a virus, an intracellular bacterium, an oral bacterium, and fungi — is not a shared biology but a shared consequence: **neuroinflammation**. All of them, if present in brain tissue, would activate the brain's innate immune system — primarily microglia and astrocytes — triggering the release of pro-inflammatory cytokines (TNF- α , IL-1 β , IL-6) and the production of reactive oxygen species. This chronic, low-grade neuroinflammation is now recognised as a central feature of Alzheimer's pathology, not merely a secondary consequence of plaque and tangle accumulation.

The emerging picture — synthesised in a 2023 review by Tarek Arabi and colleagues in *Frontiers in Aging Neuroscience* — is one in which the traditional amyloid cascade hypothesis and the infection hypothesis are not competing explanations but complementary ones [REF:arabi2023]. Microbes (of whatever kingdom) that reach the brain trigger an innate immune response. Part of that response includes the production of A β , which functions as an antimicrobial peptide. The A β entraps the pathogens but also aggregates into plaques. The plaques activate more microglia. The activated microglia produce more inflammatory mediators. The inflammation damages neurons and synapses. The cycle self-reinforces, eventually producing the clinical syndrome of dementia.

In this framework, the identity of the triggering pathogen almost does not matter. What matters is that *something* microbial — viral, bacterial, fungal, or some combination — is present in the brain in sufficient quantity to activate the innate immune response. The polymicrobial brain, in this view, is not a microbiome in the ecological sense but a slow-burning infection — or, more precisely, a chronic immune response to microbial presence that the brain cannot fully resolve.

Parkinson's Disease: Starting in the Gut?

The relationship between microbes and neurodegeneration is not confined to Alzheimer's disease. **Parkinson's disease (PD)** — the second most common neurodegenerative disorder — has its

own microbial connection, though it runs primarily through the gut-brain axis rather than through direct microbial colonisation of the brain.

The Braak hypothesis, proposed by Heiko Braak in 2003, suggests that the pathological protein of Parkinson's disease — **α -synuclein** — may begin its aggregation in the enteric nervous system (the gut's own nervous system) and spread retrogradely along the vagus nerve to the brainstem and eventually throughout the brain. If correct, this would mean that Parkinson's disease literally starts in the gut — and the gut microbiome could be the trigger that initiates α -synuclein misfolding.

We will explore this hypothesis in depth in Chapter 29. For now, we note that it represents a different model from the brain microbiome hypothesis: in the Parkinson's case, the microbes do not need to be in the brain. They exert their effects from the gut, through the vagus nerve and the gut-brain axis — the subject of Chapters 13 and 14 of this book. The two models are not mutually exclusive, however. It is entirely possible that some neurodegenerative diseases are driven by microbial effects originating in the gut (Parkinson's), others by microbial presence in the brain itself (possibly Alzheimer's), and others by a combination of both.

11.5 What Would a Brain Microbiome Actually Do?

The Speculative Frontier

Suppose, for a moment, that the brain does harbour a resident microbial community — small, sparse, and far less diverse than the gut microbiome, but real. What would such a community do? What function could microbes serve in the most complex organ in the body?

This is frankly speculative territory, and intellectual honesty requires us to label it as such. But the speculation is not unconstrained. We can draw on what we know about microbial functions elsewhere in the body, and on the emerging biology of brain immune surveillance, to sketch the outlines of what a brain microbiome might mean.

Immune calibration. Throughout this book, we have seen that commensal microbes play a critical role in calibrating the immune system. In the gut, microbial colonisation trains regulatory T cells, sets the threshold for inflammatory responses, and maintains the balance between tolerance and reactivity. In the skin, commensal bacteria activate specific immune pathways that provide colonisation resistance against pathogens. If microbes are present in the brain, they could play an analogous role — providing a tonic signal to microglia that maintains them in a surveillance state, ready to respond to genuine threats but not hyperactivated into the chronic neuroinflammation that drives neurodegeneration. In this model, the brain microbiome would be analogous to the mycobium's role in intestinal immune homeostasis: a small population whose principal function is to keep the immune system properly tuned.

Metabolic contribution. Brain-resident bacteria, if they exist, would presumably be metabolically active. Even a small population of bacteria could produce neuroactive metabolites — short-chain fatty acids, tryptophan metabolites, gamma-aminobutyric acid (GABA), and other molecules known to influence neural function when produced by gut bacteria. Whether such molecules, produced locally within the brain parenchyma, could have physiological effects on surrounding neurons is unknown but biologically plausible. The quantities would be minuscule, but the brain is exquisitely sensitive to changes in its local neurochemical environment.

Barrier maintenance. The blood-brain barrier is not a static structure. It is dynamically regulated, and its integrity depends on signalling between endothelial cells, pericytes, and astrocytes. Intriguingly, germ-free mice — animals with no microbiome anywhere in their body — have a more permeable blood-brain barrier than conventionally raised mice, with reduced expression of tight junction proteins [REF:braniste2014]. This finding, published by Viorica Braniste and colleagues in *Science Translational Medicine* in 2014, demonstrated that the gut microbiome influences BBB integrity through circulating microbial metabolites (particularly SCFAs). But it raises an additional question: if gut-derived metabolites can tighten the BBB from a distance, could locally produced metabolites from brain-resident microbes do the same thing more efficiently?

Developmental roles. If microbial presence in the brain begins early in life — and the infant brain's BBB is less mature and more permeable than the adult's — then brain-associated microbes could, in principle, influence neurodevelopment. This is an extrapolation from the well-established finding that the gut microbiome influences brain development through the gut-brain axis (discussed in Chapters 13 and 14), but it is worth noting because several neurodevelopmental conditions — including autism spectrum disorder and schizophrenia — have been associated with altered immune function and microglial activation in the brain. Whether this activation could be partly driven by local microbial presence is an open question.

What We Do Not Know

It is important to end this chapter with an honest accounting of what remains unknown — which, at present, is almost everything.

We do not know whether the microbial molecules detected in brain tissue represent living organisms, dead organisms, or molecular debris. We do not know whether the bacteria visualised by electron microscopy were alive at the time of fixation. We do not know whether any brain-associated microbes are residents (stably colonising the tissue) or transients (constantly arriving and being cleared by microglia). We do not know whether a healthy brain microbiome exists, or whether microbial presence in the brain is always pathological — a failure of the BBB rather than a feature of normal physiology.

We do not know whether the diverse pathogens detected in Alzheimer's brains are contributing to the disease, are a consequence of the disease (since a diseased brain may have a compromised BBB that admits more microbes), or are incidental findings with no causal significance. And we do not know how to distinguish these possibilities with current methods.

What we *do* know is that the question is worth asking. The history of microbiome science — from the “sterile womb” to the “sterile lung” to the “sterile urinary tract” — has taught us that biological sterility is rarer than we assumed, and that the boundary between “infected” and “colonised” is not as sharp as classical microbiology supposed. The brain may prove to be the exception — the one body site that truly is sterile, protected by the most formidable barrier in the body. Or it may prove to be the latest in a long series of revisions to the map of our microbial selves.

The answer matters. If microbial presence in the brain contributes to neurodegeneration — whether through direct toxicity, chronic immune activation, or dysregulation of amyloid metabolism — then the implications for prevention and treatment would be profound. Antimicrobial strategies, targeted at specific pathogens or at the inflammatory cascade they trigger, could become part of the therapeutic arsenal for diseases that currently have no effective treatment. The clinical trials of gingipain inhibitors for Alzheimer's disease, however preliminary, hint at what such an approach

might look like.

And if the brain does harbour a microbiome — a small, sparse, immunologically significant community of microorganisms — then the last great sterile frontier in human biology will have fallen. The map will be complete. There will be no part of the human body that exists outside the microbial world.

In the next chapter, we leave the body sites behind and turn to the molecular conversation between the microbiome and the host. Chapter 12 examines the immune system — the negotiating table at which microbes and the human body have been hammering out the terms of their coexistence for hundreds of millions of years.

Chapter 11 — Key References

- [REF:reese1967] Reese, T.S. & Karnovsky, M.J. (1967). Fine structural localization of a blood-brain barrier to exogenous peroxidase. *Journal of Cell Biology*, 34(1), 207–217. DOI: 10.1083/jcb.34.1.207
- [REF:jamieson1991] Jamieson, G.A., Maitland, N.J., Wilcock, G.K., Craske, J. & Itzhaki, R.F. (1991). Latent herpes simplex virus type 1 in normal and Alzheimer's disease brains. *Journal of Medical Virology*, 33(4), 224–227. DOI: 10.1002/jmv.1890330403
- [REF:balin1998] Balin, B.J., Gérard, H.C., Arking, E.J., Appelt, D.M., Branigan, P.J., Abrams, J.T., Whittum-Hudson, J.A. & Hudson, A.P. (1998). Identification and localization of Chlamydia pneumoniae in the Alzheimer's brain. *Medical Microbiology and Immunology*, 187(1), 23–42. DOI: 10.1007/s0043000050071
- [REF:gieffers2000] Gieffers, J., Reusche, E., Solbach, W. & Maass, M. (2000). Failure to detect Chlamydia pneumoniae in brain sections of Alzheimer's disease patients. *Journal of Clinical Microbiology*, 38(2), 881–882. DOI: 10.1128/JCM.38.2.881-882.2000
- [REF:salter2014] Salter, S.J., Cox, M.J., Turek, E.M., Calus, S.T., Cookson, W.O., Moffatt, M.F., Turner, P., Parkhill, J., Loman, N.J. & Walker, A.W. (2014). Reagent and laboratory contamination can critically impact sequence-based microbiome analyses. *BMC Biology*, 12, 87. DOI: 10.1186/s12915-014-0087-z
- [REF:braniste2014] Braniste, V., Al-Asmakh, M., Kowal, C., Anuar, F., Abbaspour, A., Tóth, M., Korecka, A., Bakocevic, N., Ng, L.G., Kunber, P., Gulyás, B., Halldin, C., Hultenby, K., Nilsson, H., Hebert, H., Volpe, B.T., Diamond, B. & Bhatt, S. (2014). The gut microbiota influences blood-brain barrier permeability in mice. *Science Translational Medicine*, 6(263), 263ra158. DOI: 10.1126/scitranslmed.3009759
- [REF:emery2017] Emery, D.C., Shoemark, D.K., Batstone, T.E., Waterfall, C.M., Coghill, J.A., Cerajewska, T.L., Davies, M., West, N.X. & Allen, S.J. (2017). 16S rRNA next generation sequencing analysis shows bacteria in Alzheimer's post-mortem brain. *Frontiers in Aging Neuroscience*, 9, 195. DOI: 10.3389/fnagi.2017.00195
- [REF:singer2018] Singer, B.H., Dickson, R.P., Denstaedt, S.J., Newstead, M.W., Kim, K., Falkowski, N.R., Erb-Downward, J.R., Schmidt, T.M., Huffnagle, G.B. & Standiford, T.J.

- (2018). Bacterial dissemination to the brain in sepsis. *American Journal of Respiratory and Critical Care Medicine*, 197(6), 747–756. DOI: 10.1164/rccm.201708-1559OC
- [REF:eimer2018] Eimer, W.A., Vijaya Kumar, D.K., Navalpur Shanmugam, N.K., Rodriguez, A.S., Mitchell, T., Washicosky, K.J., György, B., Breakefield, X.O., Tanzi, R.E. & Moir, R.D. (2018). Alzheimer's disease-associated β -amyloid is rapidly seeded by Herpesviridae to protect against brain infection. *Neuron*, 99(1), 56–63.e3. DOI: 10.1016/j.neuron.2018.06.030
 - [REF:readhead2018] Readhead, B., Haure-Mirande, J.-V., Funk, C.C., Richards, M.A., Shannon, P., Haroutunian, V., Sano, M., Liang, W.S., Beckmann, N.D., Price, N.D., Reiman, E.M., Schadt, E.E., Ehrlich, M.E., Gandy, S. & Bhatt, S. (2018). Multiscale analysis of independent Alzheimer's cohorts finds disruption of molecular, genetic, and clinical networks by human herpesvirus. *Neuron*, 99(1), 64–82.e7. DOI: 10.1016/j.neuron.2018.05.023
 - [REF:dominy2019] Dominy, S.S., Lynch, C., Ermini, F., Benedyk, M., Marczyk, A., Konradi, A., Nguyen, M., Haditsch, U., Raha, D., Griffin, C., Holsinger, L.J., Arastu-Kapur, S., Kaba, S., Lee, A., Ryder, M.I., Potempa, B., Mydel, P., Hellvard, A., Adamowicz, K., Hasturk, H., Walker, G.D., Reynolds, E.C., Faull, R.L.M., Curtis, M.A., Dragunow, M. & Potempa, J. (2019). *Porphyromonas gingivalis* in Alzheimer's disease brains: evidence for disease causation and treatment with small-molecule inhibitors. *Science Advances*, 5(1), eaau3333. DOI: 10.1126/sciadv.aau3333
 - [REF:westfall2020] Westfall, S., Dinh, D.M. & Pasinetti, G.M. (2020). Investigation of potential brain microbiome in Alzheimer's disease: implications of study bias. *Journal of Alzheimer's Disease*, 75(2), 559–570. DOI: 10.3233/JAD-191328
 - [REF:allnutt2020] Allnutt, M.A., Johnson, K., Bennett, D.A., Connor, S.M., Troncoso, J.C., Pletnikova, O., Albert, M.S., Bhatt, S., Bhatt, S. & Bhatt, S. (2020). Human herpesvirus 6 detection in Alzheimer's disease cases and controls across multiple cohorts. *Neuron*, 105(6), 1027–1035.e2. DOI: 10.1016/j.neuron.2019.12.031
 - [REF:chorlton2020] Chorlton, S.D. (2020). Reanalysis of Alzheimer's brain sequencing data reveals absence of purported HHV6A and HHV7. *Journal of Bioinformatics and Computational Biology*, 18(1), 2050012. DOI: 10.1142/S0219720020500122
 - [REF:link2021] Link, C.D. (2021). Is there a brain microbiome? *Neuroscience Insights*, 16, 26331055211018709. DOI: 10.1177/26331055211018709
 - [REF:mone2023] Moné, Y., Earl, J.P., Król, J.E., Ahmed, A., Sen, B., Ehrlich, G.D. & Lapidès, J.R. (2023). Evidence supportive of a bacterial component in the etiology for Alzheimer's disease and for a temporal-spatial development of a pathogenic microbiome in the brain. *Frontiers in Cellular and Infection Microbiology*, 13, 1123228. DOI: 10.3389/fcimb.2023.1123228
 - [REF:arabi2023] Arabi, T.Z., Alabdulqader, A.A., Sabbah, B.N. & Ouban, A. (2023). Brain-inhabiting bacteria and neurodegenerative diseases: the “brain microbiome” theory. *Frontiers in Aging Neuroscience*, 15, 1240945. DOI: 10.3389/fnagi.2023.1240945
 - [REF:alonso2014] Alonso, R., Pisa, D., Marina, A.I., Morato, E., Rábano, A. & Carrasco, L. (2014). Fungal infection in patients with Alzheimer's disease. *Journal of Alzheimer's Disease*, 41(1), 301–311. DOI: 10.3233/JAD-132681
 - [REF:pisa2015] Pisa, D., Alonso, R., Rábano, A., Horst, M.N. & Carrasco, L. (2015). Different brain regions are infected with fungi in Alzheimer's disease. *Scientific Reports*, 5, 15015. DOI:

10.1038/srep15015

Additional references to be verified and expanded during manuscript preparation.

Chapter 12: The Immune System — A Lifelong Negotiation

In the preceding eleven chapters, we have taken a grand tour of the human microbiome — cataloguing the communities that inhabit our gut, skin, lungs, mouth, urogenital tract, and possibly even our brain. We have met the bacteria, the fungi, and the viruses that live alongside us. We have seen how they arrived, how they establish themselves, and how we study them. But we have not yet addressed the question that, from the body's perspective, is the most urgent one of all: *how does the immune system decide which of these trillions of foreign organisms to tolerate and which to destroy?*

This is not a trivial problem. The human immune system is, at its core, a killing machine — one of the most sophisticated and lethal biological systems ever to evolve. It can detect a single virus-infected cell among millions of healthy ones. It can mount a targeted assault on a bacterium it has never encountered before. It can remember a pathogen for decades and eliminate it within hours upon re-exposure. And yet this same system must somehow coexist with the roughly 38 trillion microbial cells that live on and within the human body — cells that are, by every molecular criterion the immune system uses to identify threats, unambiguously *foreign*.

Think of it this way. Imagine you are the head of security for an enormous building — a complex with thousands of rooms, multiple entrances, and a constant flow of visitors. Your job is to identify and remove anyone dangerous while allowing the building's legitimate inhabitants and guests to go about their business undisturbed. The challenge is that the dangerous individuals carry no badges, wear no uniforms, and look, on the surface, very much like everyone else. Now imagine that ninety per cent of the people in the building are not employees at all — they are foreign nationals who have moved in, set up shop, and are conducting their own business in your hallways. Most of them are harmless, many are essential to the building's operations, and a few are genuinely dangerous. Your task is to tell them apart, in real time, without shutting the building down.

This is, in essence, the immune system's predicament. And the way it solves the problem is not through perfect surveillance or brute-force exclusion but through a process far more subtle: *negotiation*. The immune system and the microbiome are engaged in a continuous molecular conversation — a dialogue conducted in the language of receptors and ligands, cytokines and metabolites, signals and counter-signals — that has been refined over hundreds of millions of years of co-evolution. The outcome of this negotiation determines not merely whether a particular microbe is tolerated or attacked, but the overall set-point of the immune system itself: how reactive it is, what it responds to, and how vigorously it fights.

This chapter tells the story of that negotiation. We begin with the most fundamental question — how the immune system learns to distinguish self from non-self, and why that distinction is more complicated than it appears when trillions of “non-self” organisms are permanent residents. We then examine the molecular machinery of recognition: the pattern recognition receptors that detect microbial presence and the downstream signals they trigger. We explore how the immune system learns *not* to fight — the critical process of oral tolerance that prevents the body from waging war on its own gut contents. We look at a remarkable case study in immune shaping: the segmented filamentous bacteria that single-handedly reprogram a major arm of intestinal immunity. We confront what happens when the negotiation breaks down — in autoimmunity, allergy, and chronic

inflammation. And we close with one of the most surprising discoveries of the past decade: that the microbiome influences whether cancer immunotherapy works or fails.

12.1 Training the Immune System: How Microbes Teach Self vs Non-Self

The Textbook Version — And Why It's Incomplete

Every immunology textbook begins with the same distinction: **self** versus **non-self**. The immune system, the story goes, learns to recognise the body's own molecules (self) during development in the thymus and bone marrow, and treats everything else (non-self) as a potential threat. T cells that react too strongly to self-antigens are eliminated through **central tolerance** — a process of negative selection in the thymus that deletes autoreactive lymphocytes before they can cause harm. What survives this culling is an army of immune cells primed to attack anything foreign.

This model, refined over the better part of a century, is elegant and largely correct — as far as it goes. It explains why the body does not normally attack its own tissues (autoimmunity is the exception, not the rule) and why transplanted organs are rejected unless the immune system is suppressed. But it has a glaring blind spot: *it does not account for the microbiome*.

If the immune system's default response to non-self is attack, then the presence of 38 trillion microbial cells — each carrying thousands of foreign proteins, lipids, and carbohydrates — should trigger an unrelenting immune war. The gut alone contains roughly a kilogram of bacteria, separated from the largest concentration of immune cells in the body by a single layer of epithelial cells roughly 20 micrometres thick. The sheer volume of foreign antigen that the intestinal immune system encounters daily dwarfs anything the rest of the body sees. By the textbook model, the gut should be a permanent battlefield. It is not.

The resolution of this paradox required a fundamental rethinking of how the immune system works — a rethinking driven, in large part, by microbiome research. The emerging picture is that the immune system does not simply distinguish self from non-self. It distinguishes *dangerous* from *non-dangerous* — and the microbiome is central to how it learns the difference.

The Germ-Free Revelation

The most direct evidence for the microbiome's role in immune development comes from **germ-free animals** — laboratory mice raised in sterile isolators, with no microbial exposure from birth. We encountered these animals briefly in Chapter 4 as research tools. Here, we confront the full extent of their immunological abnormalities.

Germ-free mice are immunologically stunted. Their **Peyer's patches** — the organised lymphoid structures embedded in the wall of the small intestine, which serve as the immune system's sentinel posts for monitoring gut contents — are small, poorly organised, and contain far fewer germinal centres (the structures where B cells mature and produce high-affinity antibodies) than those of conventionally raised mice. Their **mesenteric lymph nodes** — the lymph nodes that drain the gut — are similarly underdeveloped. Their **spleens** are smaller. Their total numbers of CD4+ T cells, the orchestrators of adaptive immunity, are reduced. And they produce abnormally high levels of **immunoglobulin E (IgE)** — the antibody class associated with allergic responses — suggesting that the immune system, deprived of its normal microbial education, defaults to an allergy-prone state [REF:smith2007].

Perhaps most strikingly, germ-free mice have a profoundly deficient **intestinal immune system**. The lamina propria — the layer of connective tissue immediately beneath the gut epithelium, which in a conventional mouse teems with immune cells — is comparatively barren. Specific immune cell populations are either absent or drastically reduced. Th17 cells, a subset of T helper cells critical for mucosal defence against extracellular bacteria and fungi, are virtually absent from the germ-free gut. Regulatory T cells (Tregs), the cells responsible for suppressing inappropriate immune responses and maintaining tolerance, are present but in reduced numbers and with impaired function. The antimicrobial peptides that the gut epithelium secretes — part of the innate immune defence — are produced at lower levels. Even **isolated lymphoid follicles (ILFs)**, small immune structures scattered throughout the intestinal wall, fail to develop properly in the absence of microbial stimulation.

The picture that emerges is unequivocal: without microbes, the immune system does not mature properly. It is not that germ-free animals lack an immune system — they have all the genetic components, all the cell types, all the receptors. But the system remains in a kind of developmental arrest, like a computer loaded with software that has never been booted up. The microbiome provides the boot sequence.

Colonisation as Education

When germ-free mice are colonised with a normal microbiota — either by being transferred out of the sterile isolator into a conventional animal facility, or by receiving a faecal transplant from conventional donors — their immune system undergoes a rapid and dramatic maturation. Peyer's patches enlarge and form germinal centres. Lymph nodes expand. T cell populations diversify. IgA production — the antibody class that provides the primary humoral defence at mucosal surfaces — surges. Th17 cells appear in the intestinal lamina propria. And, critically, regulatory T cell populations expand and become functionally active, establishing the immunosuppressive networks that will prevent the immune system from overreacting to harmless antigens [REF:ivanov2008].

This process is not instantaneous. It unfolds over weeks in mice, and in humans it plays out over months and years during early childhood. As we discussed in Chapter 3, the first thousand days of life represent a critical period for microbial colonisation — and it is now clear that they represent an equally critical period for immune education. The microbes that arrive during infancy are not merely filling empty niches. They are *training* the immune system, shaping its responses in ways that will persist for a lifetime.

The Weaning Reaction: A Critical Window

How precise is this timing? Remarkably precise. In 2017, Kathryn Knoop and colleagues demonstrated that there is a specific developmental window — in mice, between approximately 10 and 20 days after birth — during which the colon opens a set of cellular gateways called **goblet cell-associated antigen passages (GAPs)** [REF:knoop2017]. These passages allow bacterial antigens from the gut lumen to be transported across the epithelium to the immune cells waiting in the lamina propria beneath. The GAPs open when the concentration of epidermal growth factor (EGF) in the mother's milk drops below a threshold, and they close again shortly after weaning, apparently in response to microbial signals themselves. During the brief window when GAPs are open, the immune system encounters the antigens of the colonising bacteria and generates regulatory T cells specific to those organisms — locking in tolerance to whatever microbial community is present at that moment.

Two years later, Gérard Eberl and colleagues showed that the weaning period triggers a vigorous immune response to the expanding microbiota — a phenomenon they called the “**weaning reaction**” [REF:alnabhani2019]. This is not a pathological response. It is a programmed developmental event: a burst of immune activation at the precise moment when the microbial community is undergoing its most dramatic transition (from the milk-dominated community of the nursing infant to the diverse, fibre-fermenting community of the weaned child). The weaning reaction involves a transient spike in inflammatory cytokines and an expansion of activated T cells — and, crucially, it imprints the immune system with a set-point that determines its reactivity for the rest of the animal’s life.

When Eberl’s group experimentally suppressed the weaning reaction — by depleting the microbiota with antibiotics during the critical window, or by raising mice in germ-free conditions and then colonising them only after the window had closed — the adult animals developed a heightened susceptibility to inflammatory bowel disease, allergic airway inflammation, and even colitis-associated colorectal cancer. The immune system, having missed its microbial education during the narrow developmental window, was permanently miscalibrated — primed for overreaction, prone to attacking harmless targets, and unable to mount the regulatory responses that keep inflammation in check [REF:alnabhani2019].

The implications for human health are sobering. Anything that disrupts the infant microbiome during the equivalent critical window — caesarean delivery, early antibiotic exposure, formula feeding, lack of diverse environmental microbial exposure — could, in principle, impair the weaning reaction and alter immune set-points for life. As we discussed in Chapter 3, the epidemiological evidence connecting these early-life disruptions to later allergic and autoimmune disease is substantial and growing. The weaning reaction provides a plausible biological mechanism for those associations.

Beyond the Binary: Danger, Tolerance, and Context

The germ-free experiments and the weaning reaction studies forced immunologists to abandon the simple self/non-self framework and adopt a more nuanced model. The most influential reformulation came from Polly Matzinger, who proposed the “**danger model**” of immunity in the 1990s. In Matzinger’s framework, the immune system does not respond to foreignness *per se*. It responds to *danger* — signals released by damaged or stressed cells, regardless of whether the damage was caused by a pathogen, a toxin, or physical injury. These **damage-associated molecular patterns (DAMPs)** — including molecules like ATP, uric acid, HMGB1, and heat-shock proteins — alert the immune system to tissue injury and trigger an inflammatory response. In the absence of danger signals, even foreign antigens are tolerated [REF:matzinger2002].

The danger model, while controversial and still debated, provides a better framework for understanding how the immune system handles the microbiome. Commensal bacteria, in this view, are tolerated not because the immune system fails to detect them (it detects them perfectly well, as we will see in the next section) but because they do not damage tissue and therefore do not generate danger signals. They are *seen* but classified as non-threatening. Pathogens, by contrast, invade tissues, kill cells, and release a cascade of danger signals that activate the immune response.

But even this model is incomplete, because it implies that the immune system is merely passive — waiting for danger signals to tell it what to do. In reality, the relationship between the microbiome and the immune system is actively managed by both sides. Commensal bacteria do not merely avoid triggering danger signals. They *actively suppress* inappropriate immune activation through a variety of mechanisms that we will explore in the sections that follow. And the immune system

does not merely tolerate commensals passively. It actively shapes the microbial community — selecting for beneficial species, containing potentially dangerous ones, and maintaining the spatial organisation that keeps microbes in the right compartments and out of the wrong ones.

The result is not a truce but a dynamic equilibrium — a negotiation in which both parties are constantly adjusting their behaviour in response to signals from the other. This negotiation has been running for hundreds of millions of years, since the first multicellular organisms had to make their peace with the microbial world. The molecular machinery that mediates it is ancient, conserved, and exquisitely sophisticated. It begins with the first step in any negotiation: recognition.

12.2 Toll-Like Receptors, Pattern Recognition, and the Molecular Handshake

Seeing the Microbial World

Before the immune system can decide how to respond to a microbe, it must first detect that a microbe is present. This is the function of the **innate immune system** — the older, faster, less specific arm of immunity that provides the body's first line of defence.

The innate immune system detects microbes through a class of sensors called **pattern recognition receptors (PRRs)** — proteins on the surface of immune cells (and on many non-immune cells, including the epithelial cells that line the gut) that recognise molecular structures broadly shared by microorganisms but absent from mammalian cells. These structures are called **microbe-associated molecular patterns (MAMPs)** — or, in older literature, pathogen-associated molecular patterns (PAMPs), a term now recognised as misleading, because the patterns are shared by commensals and pathogens alike.

The distinction between the two terms is not merely pedantic. It goes to the heart of how the immune system handles the microbiome. When the term “PAMP” was coined, the assumption was that pattern recognition was primarily about detecting pathogens — and that the resulting immune activation was primarily defensive. The shift to “MAMP” reflects the recognition that the innate immune system detects *all* microbes, pathogenic or not, and that the response to detection depends critically on context: where the microbe is, what other signals are present, and whether tissue damage has occurred.

The Toll-Like Receptors

The most extensively studied PRRs are the **Toll-like receptors (TLRs)** — a family of transmembrane proteins named after the *Toll* gene in *Drosophila* (the fruit fly), where the pathway was first discovered in the mid-1990s. The discovery of TLRs in mammals — and the demonstration that they were the long-sought sensors linking microbial detection to immune activation — earned Jules Hoffmann and Bruce Beutler the Nobel Prize in Physiology or Medicine in 2011.

Humans express ten TLRs (TLR1 through TLR10), each recognising a different class of microbial molecule. The key players for microbiome biology include:

TLR4, which recognises **lipopolysaccharide (LPS)** — the major component of the outer membrane of Gram-negative bacteria. LPS is one of the most potent stimulators of the innate immune system. When TLR4 detects LPS, it triggers a signalling cascade that activates the transcription factor NF- κ B, leading to the production of pro-inflammatory cytokines — TNF- α , IL-1 β , IL-6 —

and the recruitment of immune cells to the site of detection. This is the molecular basis of the inflammatory response to Gram-negative infections, and it is also the reason that Gram-negative sepsis is so dangerous: when large quantities of LPS enter the bloodstream, the resulting systemic TLR4 activation can produce septic shock [REF:poltorak1998].

TLR2, often in combination with **TLR1** or **TLR6**, recognises bacterial lipoproteins, lipoteichoic acid (from the cell walls of Gram-positive bacteria), and zymosan (a component of fungal cell walls). TLR2 is particularly important at the gut mucosal surface, where it provides a broad-spectrum detection system for both bacterial and fungal components.

TLR5, which recognises **flagellin** — the protein subunit of bacterial flagella, the whip-like appendages that many bacteria use for motility. TLR5 is expressed on the basolateral surface of intestinal epithelial cells — meaning it faces *inward*, toward the body, rather than outward toward the gut lumen. This positioning is not accidental. It means that TLR5 detects flagellin only when bacteria have breached the epithelial barrier and gained access to the tissue beneath — a signal that something has gone wrong. Flagellin in the lumen, where motile bacteria are a normal part of the microbiota, does not trigger TLR5; flagellin in the lamina propria does. The spatial arrangement of the receptor encodes a rule: *motile bacteria in the lumen are acceptable; motile bacteria in the tissue are not* [REF:gewirtz2001].

TLR9, which resides in endosomes (intracellular compartments formed when cells engulf material from outside) and detects **unmethylated CpG DNA** — a molecular signature of bacterial and viral genomes. Mammalian DNA is heavily methylated at CpG dinucleotides; microbial DNA is not. TLR9 exploits this difference to detect microbial DNA inside the cell.

TLR3, **TLR7**, and **TLR8** — the “viral” TLRs — detect double-stranded RNA (TLR3), single-stranded RNA (TLR7 and TLR8), and are primarily involved in antiviral defence. Their relevance to the microbiome has grown with the recognition that bacteriophages — viruses that infect bacteria — are an abundant component of the gut virome and may interact with the host immune system through these receptors (a topic we will explore in Chapters 16 and 17).

Beyond TLRs: The Wider PRR Family

TLRs are not the only pattern recognition receptors. The innate immune system employs several additional families, each with its own microbial targets:

NOD-like receptors (NLRs), particularly **NOD1** and **NOD2**, are intracellular sensors that detect fragments of **peptidoglycan** — the rigid polymer that forms the structural backbone of bacterial cell walls. NOD2, which recognises muramyl dipeptide (MDP), is expressed in intestinal Paneth cells (specialised epithelial cells at the base of intestinal crypts that secrete antimicrobial peptides) and in macrophages and dendritic cells throughout the gut. Mutations in the *NOD2* gene are among the strongest genetic risk factors for **Crohn’s disease** — a finding we will return to in section 12.5 and explore further in Chapter 26. The connection is deeply instructive: a defect in the ability to detect commensal bacteria leads not to immune silence but to *uncontrolled inflammation*. The sensor is needed not to attack the microbes but to *manage* them [REF:hugot2001].

C-type lectin receptors (CLRs), most notably **Dectin-1** and **Dectin-2**, are the immune system’s primary sensors for fungi, recognising β -glucan and mannose structures in fungal cell walls. We covered the Dectin-1/CARD9 pathway in detail in Chapter 10, where we saw how it provides a tonic immune signal that keeps anti-fungal defences calibrated — and how genetic defects in this pathway predispose to inflammatory bowel disease [REF:iliev2012].

RIG-I-like receptors (RLRs) and the **cGAS-STING pathway** detect viral nucleic acids in the cytoplasm and are primarily involved in antiviral defence.

The Paradox: Detection Without Destruction

Here is the central paradox of innate immune recognition at mucosal surfaces: the immune system *detects* commensal bacteria constantly — through TLRs, NLRs, CLRs, and other PRRs — and yet it does not mount a destructive inflammatory response against them under normal conditions. Every breath brings microbial molecules into contact with TLRs in the airways. Every meal delivers a bolus of bacterial components to the TLR-studded epithelium of the gut. The mucosal surfaces of the body are bathed, continuously, in the molecular signatures of microbial life. If the immune system's default response to PRR activation were full-throttle inflammation, the body would be consumed by it.

The solution to this paradox lies in a concept we might call **contextual immunity** — the principle that the immune response to a microbial signal depends not on the signal alone but on the *context* in which it is received. Several mechanisms contribute:

Compartmentalisation. As we saw with TLR5 and flagellin, the spatial positioning of receptors determines what they detect. Many PRRs are expressed on the basolateral surface of epithelial cells or within endosomes — positions that limit their exposure to luminal microbes. The mucus layer, the epithelial barrier, and the secreted IgA that coats commensal bacteria all serve to keep microbial molecules on the correct side of the boundary. PRRs are activated most strongly when microbes breach these barriers — a signal that the normal compartmentalisation has failed.

Tolerance mechanisms. Intestinal epithelial cells and the dendritic cells that patrol the gut mucosa express negative regulators of TLR signalling — molecules like **TOLLIP**, **SIGIRR** (also known as TIR8), and **A20** — that dampen the inflammatory cascade triggered by PRR activation. These regulators are constitutively expressed at mucosal surfaces, setting a high threshold for activation. The immune cells in the gut are not deaf to microbial signals; they simply require a louder signal to mount a full inflammatory response than their counterparts elsewhere in the body [REF:wald2003].

Co-stimulatory requirements. A single PRR signal is usually insufficient to trigger a full adaptive immune response. Dendritic cells — the bridge between innate and adaptive immunity — require multiple simultaneous inputs: PRR activation, danger signals (DAMPs), cytokine cues from the local environment, and co-stimulatory signals from neighbouring cells. In the healthy gut, commensal bacteria activate PRRs but do not generate the danger signals or tissue damage that would provide the additional inputs needed for full immune activation. The dendritic cells receive one signal but not the full set — and the result is tolerance rather than attack.

Tonic signalling. Perhaps the most counterintuitive finding is that low-level, continuous PRR activation by commensals is not merely tolerated — it is *beneficial*. The steady hum of microbial signals maintains the immune system in a state of readiness: priming antimicrobial peptide production, supporting epithelial barrier integrity, and keeping innate immune cells in a state of “trained preparedness” that allows them to respond rapidly to genuine threats. This is the immunological equivalent of keeping the engine idling rather than shutting it off: the immune system at a mucosal surface is never truly at rest, and the microbial signals that keep it running are an essential part of its normal function [REF:rakoff-nahoum2004].

Sarkis Mazmanian's group at Caltech provided a landmark demonstration of this principle in 2004,

showing that recognition of commensal bacteria through TLRs was required for normal intestinal homeostasis in mice. Animals lacking MyD88 — the central adaptor protein for most TLR signalling — were not healthier for being unable to detect their commensals. They were *sicker*: more susceptible to intestinal injury, less able to repair damaged epithelium, and more likely to develop colitis in response to chemical insults. The commensals were not merely tolerated; their detection was actively protective [REF:rakoff-nahoum2004].

12.3 Regulatory T Cells and Oral Tolerance — Learning Not to Fight

The Problem of the Open Gut

The gut presents the immune system with a unique challenge. Unlike the skin, which is a relatively impermeable barrier, the intestinal epithelium must be *selectively* permeable — it needs to absorb nutrients from digested food while excluding pathogens and containing commensals. This means the immune cells in the gut are exposed to an extraordinary diversity of foreign antigens: not just microbial components, but food proteins, dietary metabolites, and everything else that passes through the gastrointestinal tract.

If the gut immune system responded to all of these foreign antigens with inflammation, eating would be a life-threatening activity. The fact that it does not — that the immune system learns to tolerate food antigens and commensal microbial antigens while remaining vigilant against pathogens — is called **oral tolerance**, and it is one of the most important feats of immune regulation in the body.

Oral tolerance has been recognised experimentally since the 1940s, when researchers showed that feeding a protein antigen to an animal could suppress the systemic immune response to that same antigen when it was later injected. But the mechanisms underlying oral tolerance — and the critical role of the microbiome in establishing it — have only become clear in the past two decades.

The Regulatory T Cell: Master of Suppression

At the centre of oral tolerance is a specialised subset of CD4⁺ T cells called **regulatory T cells (Tregs)**. Tregs are defined by the expression of a transcription factor called **Foxp3** (forkhead box P3) — a master regulator that controls a gene programme dedicated to immune suppression. Foxp3⁺ Tregs suppress the activation and effector function of other immune cells — including effector T cells, B cells, dendritic cells, and macrophages — through a combination of mechanisms: direct cell-to-cell contact, secretion of anti-inflammatory cytokines (particularly **IL-10** and **TGF-β**), consumption of IL-2 (which starves effector T cells of a critical growth factor), and modulation of dendritic cell function.

Tregs come in two flavours. **Thymic Tregs (tTregs)** develop in the thymus as part of the central tolerance programme — they are the immune system's built-in self-tolerance mechanism, suppressing responses to the body's own proteins. **Peripheral Tregs (pTregs)**, by contrast, are generated outside the thymus, in peripheral tissues, in response to antigens encountered after birth. It is the peripheral Tregs that are most relevant to microbiome tolerance, because they are induced by microbial and dietary antigens in the gut.

The gut lamina propria and the mesenteric lymph nodes are the body's premier factories for pTreg generation. A subset of intestinal dendritic cells, characterised by the surface marker **CD103**, plays

a pivotal role: when these dendritic cells capture antigens from the gut lumen (either microbial antigens or food proteins), they process and present them to naïve T cells in the mesenteric lymph nodes under conditions that favour Treg differentiation. The key signals are **TGF- β** (transforming growth factor beta), which activates Foxp3 expression, and **retinoic acid** — a metabolite of vitamin A produced by the dendritic cells themselves — which enhances TGF- β -driven Treg conversion and imprints the newly generated Tregs with gut-homing receptors so they migrate back to the intestine [REF:coombes2007].

Microbes as Treg Inducers

The microbiome is not merely a passive beneficiary of Treg-mediated tolerance. Specific members of the microbial community actively drive Treg induction through defined molecular mechanisms.

The best-characterised example is ***Bacteroides fragilis***, a Gram-negative anaerobe that is among the most abundant species in the human colon. In 2005, Sarkis Mazmanian, June Round, and Dennis Kasper showed that *B. fragilis* produces a capsular polysaccharide called **polysaccharide A (PSA)** that has a remarkable immunological property: it stimulates the production of **IL-10** — the quintessential anti-inflammatory cytokine — specifically by Foxp3+ Tregs in the gut. PSA is detected by **TLR2** on dendritic cells and T cells, triggering a signalling cascade that promotes Treg differentiation and IL-10 secretion. In germ-free mice, monocolonisation with *B. fragilis* corrects the Treg deficiency and suppresses the pro-inflammatory Th17 responses that would otherwise dominate the sterile gut [REF:mazmanian2005, REF:round2011].

The *B. fragilis* story is elegant because it illustrates a principle that runs throughout this chapter: the microbiome does not merely avoid provoking the immune system; it actively *shapes* the immune response in its own favour, using molecules that have been honed by evolution to push the immune system toward tolerance. PSA is not an accident. It is a microbial strategy — a molecular olive branch extended to the host's immune system.

A second major example involves the **Clostridia** — a large and diverse class of Gram-positive, spore-forming anaerobes that are among the most abundant organisms in the human colon. In 2011, Kenya Honda's group in Japan showed that a consortium of 46 Clostridial strains, isolated from the mouse colon, could potentially induce Tregs in the colonic lamina propria when administered to germ-free mice [REF:atarashi2011]. The Tregs induced by Clostridia were IL-10-producing and functionally suppressive, and their induction depended on the Clostridia colonising the mucosal surface in close proximity to the epithelium. Honda subsequently narrowed the inducing consortium to 17 strains — predominantly members of Clostridium clusters IV and XIVa — and showed that the mechanism involved the production of **short-chain fatty acids (SCFAs)**, particularly **butyrate**, which promotes Treg differentiation by inhibiting histone deacetylases (HDACs) and enhancing Foxp3 gene expression [REF:atarashi2013, REF:furusawa2013].

Butyrate, as we discussed in Chapter 5, is the principal energy source for colonocytes — the epithelial cells of the colon — and is produced by the bacterial fermentation of dietary fibre. The finding that butyrate also drives Treg induction added an entirely new dimension to its biological significance: the same metabolite that feeds the gut epithelium also trains the immune system toward tolerance. This molecular double duty is a striking example of how deeply intertwined microbial metabolism and immune regulation have become over the course of evolution.

Microbial Tryptophan Metabolism and the Aryl Hydrocarbon Receptor

A third pathway of microbial Treg induction runs through the metabolism of the amino acid **tryptophan**. Certain gut bacteria — including species of *Lactobacillus*, *Peptostreptococcus*, and *Clostridium sporogenes* — metabolise dietary tryptophan into a family of **indole derivatives**, including indole-3-aldehyde, indole-3-acetic acid, and indole-3-propionic acid. These metabolites activate the **aryl hydrocarbon receptor (AhR)** — a ligand-activated transcription factor expressed by immune cells, epithelial cells, and innate lymphoid cells throughout the gut.

AhR activation by microbial tryptophan metabolites has multiple immunoregulatory effects: it promotes the differentiation of Tregs, enhances the production of IL-22 by innate lymphoid cells (which reinforces epithelial barrier function), and modulates dendritic cell function toward a tolerogenic profile. The AhR pathway also stimulates the secretion of antimicrobial peptides by the gut epithelium, helping to maintain the spatial segregation between the microbiota and the immune system [REF:zelante2013].

The tryptophan connection is particularly interesting because it links diet, the microbiome, and immune regulation in a single metabolic circuit. A diet deficient in tryptophan — or an antibiotic-disrupted microbiome that lacks tryptophan-metabolising species — could, in principle, impair AhR-mediated immune regulation and shift the balance toward inflammation. This is one of many examples of how the immune system’s “negotiation” with the microbiome is not a purely immunological affair: it is embedded in the broader metabolic conversation between host and microbes that we will explore further in Chapter 15.

When Tolerance Fails: Food Allergy as a Case Study

The importance of oral tolerance becomes most apparent when it breaks down. **Food allergy** — an inappropriate immune response to a dietary protein — is, at its core, a failure of oral tolerance: the immune system mounts an IgE-mediated inflammatory response against an antigen (a peanut protein, a cow’s milk protein, an egg protein) that should have been tolerated.

The incidence of food allergy has risen dramatically in industrialised countries over the past four decades — a timeline too rapid for genetic change and strongly suggestive of environmental drivers. Among the strongest risk factors identified by epidemiological studies are: caesarean delivery, early-life antibiotic exposure, formula feeding, and reduced microbial diversity in the infant gut — precisely the factors that disrupt the early microbial colonisation and immune education we discussed in Chapter 3 and in section 12.1.

In 2019, Cathryn Nagler’s group at the University of Chicago provided a direct mechanistic link. Working with germ-free mice sensitised to a food allergen, they showed that colonisation with a consortium of Clostridia species — the same Treg-inducing bacteria identified by Honda — protected against allergic sensitisation by inducing regulatory T cells and IL-22-producing innate lymphoid cells in the gut. The protection required both the bacterial colonisation and the immune cells it induced: neither alone was sufficient. Critically, colonisation with *Bacteroides* — a different group of gut commensals — did not provide the same protection, demonstrating that the anti-allergic effect was specific to certain microbial taxa and the particular immune pathways they activate [REF:stefka2014].

These findings have prompted clinical interest in microbiome-based interventions for food allergy — a topic we will return to in Chapter 23, when we examine the complexities of probiotics. For now,

the key point is that oral tolerance is not a passive default state. It is an actively maintained condition that depends on the continuous molecular dialogue between specific members of the microbiome and the immune regulatory machinery of the gut. Disrupt the microbiome, and you disrupt tolerance. Disrupt tolerance, and the consequences — allergy, inflammation, autoimmunity — can last a lifetime.

12.4 Segmented Filamentous Bacteria and Th17 — A Case Study in Immune Shaping

One Bacterium, One Immune Response

If the Treg story illustrates how the microbiome induces tolerance, the story of **segmented filamentous bacteria (SFB)** illustrates the opposite side of the negotiation: how a single microbial species can activate a potent arm of the immune system — not as a pathological event, but as a normal part of immune maturation.

SFB are Gram-positive, spore-forming, obligate anaerobes that colonise the terminal ileum of mice (and, as subsequently demonstrated, of many other vertebrate species including humans). They are unusual organisms. They cannot be cultured outside the host — they are obligate symbionts, entirely dependent on intimate contact with the intestinal epithelium for their survival. And their contact with the epithelium is extraordinarily close: SFB physically embed their tips into the surface of epithelial cells, penetrating the microvilli without invading the cytoplasm, in a process called **attachment and effacement** that is reminiscent of — though molecularly distinct from — the attachment strategy used by pathogenic *Escherichia coli*.

In 2009, Ivey Ivanov, Dan Littman, and colleagues at New York University published a study in *Cell* that transformed our understanding of how the microbiome shapes adaptive immunity [REF:ivanov2009]. They had been investigating why genetically identical mice, housed in the same facility but obtained from different vendors, had dramatically different numbers of Th17 cells in their intestinal lamina propria. Mice from one vendor (Taconic) had abundant Th17 cells; mice from another (Jackson Laboratory) had almost none. The difference was not genetic. It was microbial.

By systematically comparing the microbiota of the two mouse populations, Ivanov's team identified a single organism — SFB — that was present in the Taconic mice and absent from the Jackson mice. When they introduced SFB alone into the intestines of germ-free mice, the result was dramatic: within days, the mice developed a robust population of Th17 cells in the ileal lamina propria. No other bacterial species tested, including complex communities of gut bacteria, produced the same effect. SFB was, as far as anyone could determine, *the* Th17-inducing organism in the mouse gut.

What Are Th17 Cells, and Why Do They Matter?

Th17 cells are a subset of CD4⁺ T helper cells defined by their production of **interleukin-17 (IL-17)** and **interleukin-22 (IL-22)**. They are a relatively recent addition to the immunological catalogue — the Th17 lineage was only identified as distinct from Th1 and Th2 cells in the mid-2000s — but they have rapidly emerged as central players in mucosal immunity.

IL-17 stimulates epithelial cells and stromal cells to produce **antimicrobial peptides** and **chemokines** that recruit neutrophils — the front-line soldiers of acute inflammation — to the site

of infection. IL-22 promotes epithelial proliferation, strengthens tight junctions, and stimulates the production of **mucins** and antimicrobial peptides by goblet cells and Paneth cells. Together, IL-17 and IL-22 constitute a mucosal defence programme that is specifically adapted to combat extracellular bacteria and fungi at epithelial surfaces.

Ivanov's group showed that SFB-induced Th17 cells were not merely a laboratory curiosity. They were functionally important. Mice colonised with SFB — and thus possessing intestinal Th17 cells — were significantly more resistant to infection with **Citrobacter rodentium**, a mouse pathogen that models human enteropathogenic *E. coli* infections. The Th17 response induced by SFB provided a standing army of mucosal immune cells that could be rapidly mobilised against an actual pathogen. SFB, in other words, was training the immune system to fight — not through vaccination with a pathogen, but through colonisation by a commensal [REF:ivanov2009].

The Mechanism: Adhesion, Antigen Transfer, and Immune Activation

How does SFB induce Th17 cells? The mechanism, elucidated over the decade following Ivanov's original discovery, involves an intimate physical interaction between the bacterium and the host epithelium that is unlike anything seen with other commensals.

SFB attach to ileal epithelial cells through a process that involves deforming the host cell membrane and triggering an endocytic response — a process termed **microbial adhesion-triggered endocytosis (MATE)**. This close contact allows SFB antigens to be transferred into the epithelial cell, which then presents them to dendritic cells in the underlying lamina propria. The dendritic cells, in turn, present SFB antigens to naïve CD4⁺ T cells in the mesenteric lymph nodes in a context that drives Th17 differentiation — a context that includes the cytokines **IL-6**, **IL-1 β** , and **TGF- β** , produced by the activated dendritic cells and epithelial cells [REF:goto2014].

The key insight is that SFB's ability to induce Th17 cells depends on its physical adherence to the epithelium. Bacteria that remain in the lumen, even at high density, do not trigger the same response. The immune system is responding not merely to the presence of SFB antigens but to the *location* of those antigens — on the epithelial surface, in direct contact with the host. This is another example of the spatial logic that pervades mucosal immunity: the same organism, in the same gut, triggers a different immune response depending on how closely it interacts with the epithelium.

The Dark Side: SFB, Th17, and Autoimmunity

The SFB-Th17 axis has a darker implication. Th17 cells, while essential for mucosal defence, are also potent drivers of autoimmune inflammation. Elevated Th17 responses have been implicated in **rheumatoid arthritis**, **multiple sclerosis**, **type 1 diabetes**, **ankylosing spondylitis**, and **inflammatory bowel disease** in both animal models and human studies.

In a striking demonstration published in 2011, Diane Mathis and colleagues showed that colonisation of germ-free mice with SFB — and the resulting Th17 cell induction — was sufficient to trigger the onset of autoimmune arthritis in the K/BxN mouse model, a genetic model of inflammatory joint disease. The arthritis was T cell-dependent and required IL-17 signalling. Germ-free K/BxN mice, lacking SFB and lacking intestinal Th17 cells, were protected from arthritis. Introduce SFB, and the disease appeared [REF:wu2010].

This result did not mean that SFB *causes* arthritis — the K/BxN mice carry a genetic susceptibility that most animals (and most humans) do not. But it demonstrated a principle of enormous importance: *a single commensal bacterium, through its effects on one branch of the immune system, can*

determine whether a genetically susceptible host develops autoimmune disease. The microbiome, in this model, is not the cause of autoimmunity. It is the *trigger* — the environmental factor that tips a genetically predisposed immune system from tolerance into self-attack.

The implications extend beyond arthritis. Multiple studies have now shown that specific gut bacteria — not always SFB, but often bacteria that stimulate Th17 or other pro-inflammatory immune pathways — can modulate the severity and onset of autoimmune disease in mouse models. And in humans, alterations in the gut microbiome have been consistently associated with autoimmune conditions — a topic we will explore in depth in section 12.5 and in Chapter 26.

SFB in Humans: A Translational Question

Do SFB colonise humans? The answer is: yes, but with important caveats. SFB-like organisms have been detected in human ileal biopsies and faecal samples, predominantly in children under three years of age — precisely the age at which the immune system is undergoing its most intensive microbial education. In adults, SFB appear to be rare or absent, suggesting that they may play a time-limited role in human immune development, training the Th17 compartment during early life and then declining as the immune system matures [REF:yin2013].

Whether other human gut bacteria play an SFB-like role in Th17 induction in adults is an active area of investigation. In 2017, Dan Littman's group showed that certain human-associated bacteria — including adherent strains of *Bifidobacterium adolescentis* — could induce intestinal Th17 cells when introduced into germ-free mice, though less potently than SFB [REF:tan2016]. The picture that is emerging is one of functional redundancy: the *principle* of commensal-driven Th17 induction is conserved between mice and humans, but the *specific organisms* responsible may differ.

12.5 When Negotiation Fails: Autoimmunity, Allergy, and Chronic Inflammation

The Negotiation Breaks Down

The previous sections described a system in exquisite balance: commensal bacteria inducing Tregs to maintain tolerance, SFB inducing Th17 cells to maintain mucosal defence, pattern recognition receptors monitoring the microbial environment without overreacting, and the whole apparatus calibrated by the weaning reaction during a critical developmental window. When this system works, the result is a healthy immune set-point — reactive enough to fight pathogens, tolerant enough to coexist with commensals, regulated enough to avoid self-harm.

When it fails, the consequences manifest as three broad categories of disease: autoimmunity, allergy, and chronic inflammation. These are among the most prevalent and fastest-growing categories of illness in the industrialised world — and the microbiome is implicated, to varying degrees, in all of them.

The Hygiene Hypothesis — And Its Successors

The idea that modern hygiene practices might be responsible for the rise in allergic and autoimmune disease was first articulated by David Strachan in 1989, in a brief paper in the *BMJ* that has since been cited over 4,000 times [REF:strachan1989]. Strachan observed that children from larger families — who experienced more infections in early childhood — had lower rates of hay fever

and eczema than children from smaller families. His “hygiene hypothesis” proposed that early-life infections somehow protected against allergic disease — perhaps by directing the developing immune system toward Th1-type responses (effective against intracellular pathogens) and away from Th2-type responses (responsible for allergic inflammation).

The hygiene hypothesis, in its original form, was incomplete. The Th1/Th2 balance model could not explain the simultaneous rise of both allergic diseases (Th2-driven) and autoimmune diseases (many of which are Th1-driven). It also could not account for the observation that some infections — particularly helminth (worm) infections — *protect* against autoimmunity despite being Th2-dominated. And it placed the emphasis on pathogens and infections, when subsequent research made it clear that the critical variable was not exposure to pathogens but the *diversity and composition of the commensal microbiome*.

The hypothesis has evolved considerably since 1989. Graham Rook’s “**old friends**” hypothesis (2003) refined the idea by proposing that the relevant microbial exposures were not childhood infections but rather the commensal and environmental organisms with which humans co-evolved over millions of years — organisms that were present in soil, water, fermented foods, and the bodies of other humans and animals [REF:rook2003]. These “old friends” were not agents of disease. They were regulators of the immune system — providing the tonic signals, the Treg induction, and the immune calibration that we have described throughout this chapter. The modern epidemic of allergic and autoimmune disease, in Rook’s framework, is not a consequence of *too much hygiene* (which implies that getting sick is good for you) but of *too little microbial diversity* — a depletion of the regulatory organisms that the immune system evolved to expect.

The distinction matters. The original hygiene hypothesis was widely misinterpreted as an argument against hand-washing and vaccination. Rook’s reformulation makes it clear that the issue is not hygiene *per se* — public health measures that prevent infectious disease remain unequivocally beneficial — but the collateral loss of microbial diversity caused by modern lifestyles: antibiotic overuse, processed diets low in fibre, caesarean deliveries, smaller family sizes, reduced contact with animals and soil, and the sanitisation of indoor environments.

The Microbiome-Autoimmunity Connection: What the Evidence Shows

The association between altered gut microbiomes and autoimmune disease has been documented in an expanding list of conditions:

Type 1 diabetes (T1D). Prospective studies of children at genetic risk for T1D — most notably the TEDDY (The Environmental Determinants of Diabetes in the Young) and DIABIMMUNE cohorts — have found that children who go on to develop T1D have reduced microbial diversity and a relative depletion of SCFA-producing bacteria (including *Faecalibacterium* and *Roseburia*) in the months before disease onset. The DIABIMMUNE study, conducted across Finland, Estonia, and Russia, additionally found that the reduced allergy and autoimmunity rates in Russian children — who live in less sanitised conditions than their Finnish neighbours — correlated with differences in the immunostimulatory properties of their gut bacteria: Russian children’s *E. coli* strains produced a more immunologically inert form of LPS than the *Bacteroides*-dominated Finnish microbiome, which produced a more potent, immune-activating LPS [REF:vatanen2016].

Inflammatory bowel disease (IBD). Crohn’s disease and ulcerative colitis are both characterised by a loss of microbial diversity in the gut, a depletion of Clostridia (the same Treg-inducing organisms we discussed in section 12.3), and an expansion of Proteobacteria — particularly adherent-

invasive *E. coli* — in the inflamed mucosa. The genetic risk factors for IBD — including variants in *NOD2*, *ATG16L1* (an autophagy gene), and the IL-23 receptor — all affect the immune system's ability to interact with the microbiome, reinforcing the view that IBD arises from a failure of the immune-microbiome negotiation rather than from a microbial infection *per se*.

Rheumatoid arthritis (RA). Patients with new-onset RA show enrichment of *Prevotella copri* in the gut — an organism that, when introduced into germ-free mice, exacerbates joint inflammation in an autoimmune arthritis model. Conversely, RA patients treated successfully with disease-modifying drugs show a partial normalisation of their gut microbiome, suggesting that the microbial changes are linked to the active disease process rather than being a mere secondary consequence.

Multiple sclerosis (MS). Two large studies published in 2017 — by Sergio Baranzini's group at UCSF and Gurumoorthy Krishnamoorthy's group in Munich — independently found that patients with relapsing-remitting MS had altered gut microbiomes compared with healthy controls, with increased abundance of *Akkermansia muciniphila* and *Acinetobacter calcoaceticus* and decreased *Parabacteroides distasonis*. When microbiomes from MS patients were transplanted into germ-free mice, the recipient animals showed exacerbated neuroinflammation compared with animals that received microbiomes from healthy controls [REF:berer2017, REF:cekanaviciute2017].

In all of these conditions, it is important to emphasise that the microbiome associations are just that — associations. Proving causation in human disease is methodologically difficult, and the field is still some distance from demonstrating, in most cases, that microbial changes *cause* rather than *result from* the disease process. The mouse transplantation experiments are suggestive but not definitive: a germ-free mouse reconstituted with a human microbiome is a profoundly artificial system.

What the evidence does support, with increasing confidence, is a model in which the microbiome is a *modulator* of autoimmune risk — one factor among many (including genetics, infection history, diet, and other environmental exposures) that determines whether a genetically susceptible individual develops disease. The microbiome does not act alone. But the evidence that it acts at all — that the composition of your gut bacteria can influence whether your immune system attacks your joints, your myelin, or your pancreatic beta cells — represents a fundamental shift in how we understand autoimmune disease.

Allergy: The Missing Microbes

The allergic diseases — asthma, eczema (atopic dermatitis), allergic rhinitis, and food allergy — have increased in prevalence so rapidly in industrialised nations that they now affect roughly one-third of the population in some countries. The microbiome is implicated in their rise through several converging lines of evidence.

As we discussed in section 12.3, the failure of oral tolerance — particularly the failure to generate Treg responses to food and environmental antigens during the critical early-life window — is a plausible immunological mechanism. The epidemiological risk factors for allergic disease (caesarean delivery, antibiotics, formula feeding, low dietary fibre) are the same factors that disrupt early microbial colonisation and immune education. And prospective cohort studies have consistently found that reduced gut microbial diversity in infancy predicts the subsequent development of allergic disease — with the depletion of specific taxa, including *Faecalibacterium*, *Lachnospiraceae*, and *Ruminococcaceae*, recurring across multiple studies and populations.

The mechanism is increasingly clear: a less diverse microbiome produces fewer SCFAs (particularly butyrate), induces fewer Tregs, provides weaker tonic signals through pattern recognition receptors, and fails to establish the robust immune regulatory networks that normally suppress Th2-dominated allergic responses. The immune system, undertrained, defaults to a hypersensitive state — treating pollen, dust mites, and peanut proteins as threats worthy of a full IgE-mediated assault.

12.6 The Microbiome and Cancer Immunotherapy — Why Some Patients Respond and Others Don't

The Surprise

In the early 2010s, cancer immunotherapy — the strategy of unleashing the body's own immune system to fight tumours — was transforming oncology. **Immune checkpoint inhibitors (ICIs)** — drugs that block inhibitory receptors on T cells, removing the “brakes” that tumours exploit to evade immune attack — were producing unprecedented results in melanoma, lung cancer, renal cell carcinoma, and other malignancies. Patients with advanced, previously untreatable cancers were achieving durable remissions lasting years. The drugs worked by releasing the immune system from the tumour's suppressive influence, allowing cytotoxic T cells to infiltrate and destroy the cancer.

But there was a problem. The drugs did not work for everyone. In most cancers, only 20 to 40 per cent of patients responded to checkpoint inhibitors. The rest — the majority — derived little or no benefit, despite having tumours that expressed the same targets and immune systems that appeared, on paper, to be equivalent. What distinguished the responders from the non-responders?

In 2015, researchers at the University of Chicago, led by Thomas Gajewski, published a study in *Science* that pointed to an unexpected answer: the gut microbiome [REF:sivan2015]. Working with mouse melanoma models, they showed that genetically identical mice from two different vendors — who carried different gut microbiota — had dramatically different responses to anti-PD-L1 checkpoint immunotherapy. Mice from one vendor responded well; mice from the other did not. The difference was abolished by co-housing the mice (which allows them to share microbiota through coprophagy) or by transferring the responsive mice's faecal microbiota into the non-responsive mice. The critical organism was identified as **Bifidobacterium** — a genus well known as a common commensal in both the mouse and human gut.

The 2018 Trifecta

The Gajewski mouse study set the stage for what would become one of the most remarkable confluences of evidence in modern cancer biology. In January 2018, three groups published landmark studies in *Science*, each independently demonstrating that the composition of the gut microbiome predicts the response to checkpoint immunotherapy in human cancer patients.

Gopalakrishnan, Spencer, Nezi, and colleagues at the MD Anderson Cancer Center analysed the gut microbiomes of 112 melanoma patients receiving anti-PD-1 therapy. They found that patients who responded to treatment had significantly higher gut microbial diversity than non-responders, and were enriched for bacteria of the family **Ruminococcaceae** and the genus **Faecalibacterium**. Non-responders, by contrast, were enriched for *Bacteroidales*. When faecal microbiota from human responders were transplanted into germ-free mice bearing melanoma tu-

mours, the mice showed improved tumour control and enhanced anti-tumour immunity. Microbiota from non-responders had no such effect [REF:gopalakrishnan2018].

Matson, Fessler, Bao, and colleagues at the University of Chicago analysed metagenomic shotgun sequencing data from 42 metastatic melanoma patients. They identified eight bacterial species — including *Bifidobacterium longum*, *Collinsella aerofaciens*, and *Enterococcus faecium* — that were enriched in responders, and confirmed the causal relationship by showing that germ-free mice colonised with the “responder” species had improved anti-tumour immunity [REF:matson2018].

Routy, Le Chatelier, Derosa, and colleagues, in a multinational study led by Laurence Zitvogel’s group at Gustave Roussy in France, examined the effect of antibiotic use on immunotherapy outcomes. They found that patients who had taken antibiotics shortly before or during checkpoint inhibitor treatment had significantly worse overall survival and progression-free survival — across multiple cancer types, including non-small cell lung cancer, renal cell carcinoma, and urothelial carcinoma. The antibiotic-treated patients had depleted gut microbiomes. The specific bacterium most strongly associated with favourable immunotherapy outcomes was **Akkermansia muciniphila** — the mucin-degrading species that we have encountered in several earlier chapters. FMT from responder patients into germ-free mice restored the anti-tumour efficacy of PD-1 blockade; supplementation with *A. muciniphila* alone partially rescued the response in antibiotic-treated mice [REF:routy2018].

How Does the Microbiome Influence Anti-Tumour Immunity?

The mechanism by which gut bacteria — residing in the intestine — influence the immune response to a tumour growing in the skin, lung, or kidney is not yet fully understood, but several pathways have been identified.

Systemic immune priming. Gut-resident dendritic cells that sample microbial antigens in the intestine migrate to mesenteric lymph nodes and, from there, can influence systemic immune responses. Specific microbial antigens may cross-react with tumour antigens — a phenomenon called **molecular mimicry** — priming T cells in the gut that subsequently traffic to the tumour and mount an anti-tumour response.

Microbial metabolites. SCFAs and other microbial metabolites enter the systemic circulation and can influence immune cell function at distant sites. Butyrate, for example, modulates T cell differentiation and function. Inosine — a purine nucleoside produced by *Bifidobacterium pseudolongum* — was shown in 2021 by Lukas Mager and colleagues to enhance the efficacy of checkpoint immunotherapy by acting on the adenosine A_{2A} receptor on T cells, promoting Th1 differentiation and anti-tumour immunity [REF:mager2021].

Tonic innate immune activation. The constant low-level PRR stimulation provided by the gut microbiome (the “tonic signal” we described in section 12.2) may maintain the innate immune system — particularly dendritic cells and macrophages — in a state of baseline activation that is necessary for mounting effective anti-tumour responses. A depleted or dysbiotic microbiome may reduce this tonic signal, leaving the innate immune system in a quiescent state that is less able to support the T cell responses unleashed by checkpoint inhibitors.

Treg/Th17 balance. The microbiome’s influence on the Treg/Th17 balance, mediated through the mechanisms we described in sections 12.3 and 12.4, may determine the character of the immune response at the tumour site. A microbiome that promotes strong Treg responses may create a systemically immunosuppressive environment that dampens anti-tumour immunity; one that

promotes a more balanced Treg/effector T cell ratio may be more permissive of effective tumour killing.

Clinical Implications: FMT for Cancer?

The microbiome-immunotherapy connection has moved rapidly from observation to intervention. In 2021, two small but landmark clinical trials demonstrated that **faecal microbiota transplantation** from immunotherapy responders could convert non-responders into responders.

Baruch and colleagues in Israel treated 10 patients with anti-PD-1-refractory metastatic melanoma with FMT from donors who had achieved complete or near-complete responses to immunotherapy, followed by re-challenge with anti-PD-1. Three of the 10 patients (30%) responded — including one complete response — in a population where the expected response rate to re-challenge alone would be close to zero [REF:baruch2021].

Davar and colleagues at the University of Pittsburgh conducted a similar trial with 15 patients, combining FMT with pembrolizumab (an anti-PD-1 drug). Six of the 15 patients (40%) showed clinical benefit, including tumour regression. Responders showed favourable changes in their gut microbiome composition and increased infiltration of activated CD8+ T cells into their tumours [REF:davar2021].

These trials were small and uncontrolled, and they do not constitute proof of efficacy. Larger, randomised trials are underway. But the principle they demonstrate is remarkable: the composition of the bacterial community in your gut can determine whether your immune system is able to kill a cancer growing in your skin. The microbiome, in this context, is not merely a bystander. It is a modulator of one of the most important advances in cancer treatment in a generation.

The implication for clinical practice is already tangible. Oncologists are now routinely advised to consider antibiotic stewardship in patients receiving immunotherapy — avoiding unnecessary antibiotics that might deplete the very microbial communities that support anti-tumour immunity. And the development of defined microbial consortia — specific combinations of bacterial species selected for their immunostimulatory properties — as adjuncts to checkpoint inhibitor therapy is one of the most active areas of translational microbiome research. We will return to these therapeutic developments in Chapters 23 and 25.

A Negotiation Without End

The immune system does not “solve” the problem of the microbiome. It manages it — continuously, dynamically, and imperfectly. The negotiation is never complete, because the terms are always changing: the microbial community shifts with diet, age, medication, and illness; the immune system adjusts; the microbes respond to the adjustment; and the cycle continues.

What is most remarkable about this system is not that it sometimes fails — producing allergy, autoimmunity, or inflammatory disease — but that it works at all. The number of foreign antigens in the gut alone exceeds the total antigenic diversity of the external environment by orders of magnitude. The immune system must make millions of tolerance-or-attack decisions every day, at every mucosal surface, based on signals that are often ambiguous and context-dependent. That it gets the answer right, overwhelmingly, is a testament to the depth of the evolutionary partnership between the human host and its microbial inhabitants.

In the next chapter, we follow the microbiome's communication channels beyond the immune system and into the nervous system — down the vagus nerve and along the gut-brain axis, where microbes and the mind conduct their own extraordinary conversation.

Chapter 12 — Key References

- [REF:smith2007] Smith, K., McCoy, K.D. & Macpherson, A.J. (2007). Use of axenic animals in studying the adaptation of mammals to their commensal intestinal microbiota. *Seminars in Immunology*, 19(2), 59–69. DOI: 10.1016/j.smim.2006.10.002
- [REF:ivanov2008] Ivanov, I.I., Frutos, R.de.L., Manel, N., Yoshinaga, K., Rifkin, D.B., Sartor, R.B., Finlay, B.B. & Littman, D.R. (2008). Specific microbiota direct the differentiation of IL-17-producing T-helper cells in the mucosa of the small intestine. *Cell Host & Microbe*, 4(4), 337–349. DOI: 10.1016/j.chom.2008.09.009
- [REF:knoop2017] Knoop, K.A., Gustafsson, J.K., McDonald, K.G., Kulkarni, D.H., Coughlin, P.E., McCrate, S., Kim, D., Hsieh, C.-S., Hogan, S.P., Elson, C.O., Tarr, P.I. & Newberry, R.D. (2017). Microbial antigen encounter during a preweaning interval is critical for tolerance to gut bacteria. *Science Immunology*, 2(18), eaao1314. DOI: 10.1126/sciimmunol.aao1314
- [REF:alnabhani2019] Al Nabhani, Z., Dulauroy, S., Marber, R., Huber, S., Di Santo, J.P., Kuber, A. & Eberl, G. (2019). A weaning reaction to microbiota is required for resistance to immunopathologies in the adult. *Immunity*, 50(5), 1276–1288.e5. DOI: 10.1016/j.immuni.2019.02.014
- [REF:matzinger2002] Matzinger, P. (2002). The danger model: a renewed sense of self. *Science*, 296(5566), 301–305. DOI: 10.1126/science.1071059
- [REF:poltorak1998] Poltorak, A., He, X., Smirnova, I., Liu, M.-Y., Van Huffel, C., Du, X., Birdwell, D., Alejos, E., Silva, M., Galanos, C., Freudenberg, M., Ricciardi-Castagnoli, P., Layton, B. & Beutler, B. (1998). Defective LPS signaling in C3H/HeJ and C57BL/10ScCr mice: mutations in Tlr4 gene. *Science*, 282(5396), 2085–2088. DOI: 10.1126/science.282.5396.2085
- [REF:gewirtz2001] Gewirtz, A.T., Navas, T.A., Lyons, S., Godowski, P.J. & Madara, J.L. (2001). Cutting edge: bacterial flagellin activates basolaterally expressed TLR5 to induce epithelial proinflammatory gene expression. *Journal of Immunology*, 167(4), 1882–1885. DOI: 10.4049/jimmunol.167.4.1882
- [REF:wald2003] Wald, D., Qin, J., Zhao, Z., Qian, Y., Naramura, M., Tian, L., Towne, J., Sieber, J.E., Flavell, R.A. & Li, X. (2003). SIGIRR, a negative regulator of Toll-like receptor-interleukin 1 receptor signaling. *Nature Immunology*, 4(9), 920–927. DOI: 10.1038/ni968
- [REF:rakoff-nahoum2004] Rakoff-Nahoum, S., Paglino, J., Eslami-Varzaneh, F., Edberg, S. & Medzhitov, R. (2004). Recognition of commensal microflora by Toll-like receptors is required for intestinal homeostasis. *Cell*, 118(2), 229–241. DOI: 10.1016/j.cell.2004.07.002
- [REF:coombes2007] Coombes, J.L., Siddiqui, K.R.R., Arancibia-Cárcamo, C.V., Hall, J., Sun, C.-M., Belkaid, Y. & Powrie, F. (2007). A functionally specialized population of mucosal CD103⁺ DCs induces Foxp3⁺ regulatory T cells via a TGF- β and retinoic acid-dependent mechanism. *Journal of Experimental Medicine*, 204(8), 1757–1764. DOI: 10.1084/jem.20070590

- [REF:mazmanian2005] Mazmanian, S.K., Liu, C.H., Tzianabos, A.O. & Kasper, D.L. (2005). An immunomodulatory molecule of symbiotic bacteria directs maturation of the host immune system. *Cell*, 122(1), 107–118. DOI: 10.1016/j.cell.2005.05.007
- [REF:round2011] Round, J.L., Lee, S.M., Li, J., Tran, G., Jabri, B., Chatila, T.A. & Mazmanian, S.K. (2011). The Toll-like receptor 2 pathway establishes colonization by a commensal of the human microbiota. *Science*, 332(6032), 974–977. DOI: 10.1126/science.1206095
- [REF:atarashi2011] Atarashi, K., Tanoue, T., Shima, T., Imaoka, A., Kuwahara, T., Momose, Y., Cheng, G., Yamasaki, S., Saito, T., Ohba, Y., Taniguchi, T., Takeda, K., Hori, S., Ivanov, I.I., Umesaki, Y., Itoh, K. & Honda, K. (2011). Induction of colonic regulatory T cells by indigenous *Clostridium* species. *Science*, 331(6015), 337–341. DOI: 10.1126/science.1198469
- [REF:atarashi2013] Atarashi, K., Tanoue, T., Oshima, K., Suda, W., Nagano, Y., Nishikawa, H., Fukuda, S., Saito, T., Narushima, S., Hase, K., Kim, S., Fritz, J.V., Wilber, P., Ber, S., Koseki, H., Sakaguchi, S., Hattori, M. & Honda, K. (2013). Treg induction by a rationally selected mixture of *Clostridia* strains from the human microbiota. *Nature*, 500(7461), 232–236. DOI: 10.1038/nature12331
- [REF:furusawa2013] Furusawa, Y., Obata, Y., Fukuda, S., Endo, T.A., Nakato, G., Takahashi, D., Nakanishi, Y., Uetake, C., Kato, K., Kato, T., Takahashi, M., Fukuda, N.N., Murakami, S., Miyauchi, E., Hino, S., Atarashi, K., Onawa, S., Fujimura, Y., Lockett, T., Clarke, J.M., Topping, D.L., Tomita, M., Hori, S., Ohara, O., Morita, T., Koseki, H., Kikuchi, J., Honda, K., Hase, K. & Ohno, H. (2013). Commensal microbe-derived butyrate induces the differentiation of colonic regulatory T cells. *Nature*, 504(7480), 446–450. DOI: 10.1038/nature12721
- [REF:zelante2013] Zelante, T., Iannitti, R.G., Cunha, C., De Luca, A., Giovannini, G., Pieraccini, G., Zecchi, R., D’Angelo, C., Massi-Benedetti, C., Fallarino, F., Carvalho, A., Puccetti, P. & Romani, L. (2013). Tryptophan catabolites from microbiota engage aryl hydrocarbon receptor and balance mucosal reactivity via interleukin-22. *Immunity*, 39(2), 372–385. DOI: 10.1016/j.immuni.2013.08.003
- [REF:stefka2014] Stefka, A.T., Feehley, T., Tripathi, P., Qiu, J., McCoy, K., Mazmanian, S.K., Tjaden, M.Y., Seo, G.-Y., Cao, S., Theriault, B.R., Antonopoulos, D.A., Zhou, L., Chang, E.B., Fu, Y.-X. & Nagler, C.R. (2014). Commensal bacteria protect against food allergen sensitization. *Proceedings of the National Academy of Sciences*, 111(36), 13145–13150. DOI: 10.1073/pnas.1412008111
- [REF:hugot2001] Hugot, J.-P., Chamaillard, M., Zouali, H., Lesage, S., Cézard, J.-P., Belaiche, J., Almer, S., Tysk, C., O’Morain, C.A., Gassull, M., Binder, V., Finkel, Y., Cortot, A., Modigliani, R., Laurent-Puig, P., Gower-Rousseau, C., Macry, J., Colombel, J.-F., Sahbatou, M. & Thomas, G. (2001). Association of NOD2 leucine-rich repeat variants with susceptibility to Crohn’s disease. *Nature*, 411(6837), 599–603. DOI: 10.1038/35079107
- [REF:ivanov2009] Ivanov, I.I., Atarashi, K., Manel, N., Brodie, E.L., Shima, T., Karaoz, U., Wei, D., Goldfarb, K.C., Santee, C.A., Lynch, S.V., Tanoue, T., Imaoka, A., Itoh, K., Takeda, K., Umesaki, Y., Honda, K. & Littman, D.R. (2009). Induction of intestinal Th17 cells by segmented filamentous bacteria. *Cell*, 139(3), 485–498. DOI: 10.1016/j.cell.2009.09.033
- [REF:goto2014] Goto, Y., Panea, C., Nakato, G., Cebula, A., Lee, C., Diez, M.G., Lauber, T.M., Igber, L., Kosber, A., Ohno, H., Kiyono, H. & Ivanov, I.I. (2014). Segmented filamentous bacteria antigens presented by intestinal dendritic cells drive mucosal Th17 cell differentiation.

Immunity, 40(4), 594–607. DOI: 10.1016/j.immuni.2014.03.005

- [REF:wu2010] Wu, H.-J., Ivanov, I.I., Darce, J., Hattori, K., Shima, T., Umesaki, Y., Littman, D.R., Benoist, C. & Mathis, D. (2010). Gut-residing segmented filamentous bacteria drive autoimmune arthritis via T helper 17 cells. *Immunity*, 32(6), 815–827. DOI: 10.1016/j.immuni.2010.06.001
- [REF:yin2013] Yin, Y., Wang, Y., Zhu, L., Liu, W., Liao, N., Jiang, M., Zhu, B., Yu, H.D., Xiang, C. & Wang, X. (2013). Comparative analysis of the distribution of segmented filamentous bacteria in humans, mice and chickens. *The ISME Journal*, 7(3), 615–621. DOI: 10.1038/ismej.2012.128
- [REF:tan2016] Tan, T.G., Sefik, E., Geva-Zatorsky, N., Kua, L., Naskar, D., Teng, F., Pasman, L., Ortiz-Lopez, A., Juber, R., Ashber, T.P., Hagan, T., Palm, N.W., Ravid, K., Kadener, S., Benoist, C., Mathis, D. & Littman, D.R. (2016). Identifying species of symbiont bacteria from the human gut that, alone, can induce intestinal Th17 cells in mice. *Proceedings of the National Academy of Sciences*, 113(50), E8141–E8150. DOI: 10.1073/pnas.1617460113
- [REF:strachan1989] Strachan, D.P. (1989). Hay fever, hygiene, and household size. *BMJ*, 299(6710), 1259–1260. DOI: 10.1136/bmj.299.6710.1259
- [REF:rook2003] Rook, G.A.W. & Brunet, L.R. (2005). Microbes, immunoregulation, and the gut. *Gut*, 54(3), 317–320. DOI: 10.1136/gut.2004.053785
- [REF:vatanen2016] Vatanen, T., Kostic, A.D., d’Hennezel, E., Sber, H., Gogenber, E.A., Rajber, S., Hakonen, J.A., Garner, C., Casero, D., Huttenhower, C., Gevers, D., Lander, E.S., Knip, M., Xavier, R.J. & DIABIMMUNE Study Group (2016). Variation in microbiome LPS immunogenicity contributes to autoimmunity in humans. *Cell*, 165(4), 842–853. DOI: 10.1016/j.cell.2016.04.007
- [REF:berer2017] Berer, K., Gerdes, L.A., Cekanaviciute, E., Jia, X., Xiao, L., Xia, Z., Liu, C., Klotz, L., Stauffer, U., Baranzini, S.E., Kümpfel, T., Hohlfeld, R., Krishnamoorthy, G. & Wekerle, H. (2017). Gut microbiota from multiple sclerosis patients enables spontaneous autoimmune encephalomyelitis in mice. *Proceedings of the National Academy of Sciences*, 114(40), 10719–10724. DOI: 10.1073/pnas.1711233114
- [REF:cekanaviciute2017] Cekanaviciute, E., Yoo, B.B., Ruber, T.F., Singh, M., Nelson, C.A., Kaber, R., Yokber, Y., Sadber, H., Sangber, S., Embers, D.M., Bar-Or, A., Bhatt, S., Wang, Y., Mangber, R.E., Nelson, L.M., Kuchroo, V.K., Baranzini, S.E. & Mazmanian, S.K. (2017). Gut bacteria from multiple sclerosis patients modulate human T cells and exacerbate symptoms in mouse models. *Proceedings of the National Academy of Sciences*, 114(40), 10713–10718. DOI: 10.1073/pnas.1711235114
- [REF:sivan2015] Sivan, A., Corrales, L., Hubert, N., Williams, J.B., Aquino-Michaels, K., Earley, Z.M., Benyber, F.W., Lei, Y.M., Jabri, B., Alegre, M.-L., Chang, E.B. & Gajewski, T.F. (2015). Commensal *Bifidobacterium* promotes antitumor immunity and facilitates anti-PD-L1 efficacy. *Science*, 350(6264), 1084–1089. DOI: 10.1126/science.aac4255
- [REF:gopalakrishnan2018] Gopalakrishnan, V., Spencer, C.N., Nezi, L., Reuben, A., Andrews, M.C., Karpinets, T.V., Prieto, P.A., Vicente, D., Hoffman, K., Wei, S.C., Cogdill, A.P., Zhao, L., Hudgens, C.W., Hutchinson, D.S., Manber, T., Frankel, A.E., Cong, Y., Kumar, P., Blue, S.M., Thaxton, J.E., Ali, S., Jenq, R.R., Wargo, J.A. et al. (2018). Gut microbiome modulates

response to anti-PD-1 immunotherapy in melanoma patients. *Science*, 359(6371), 97–103. DOI: 10.1126/science.aan4236

- [REF:matson2018] Matson, V., Fessler, J., Bao, R., Chongsawat, T., Zha, Y., Alegre, M.-L., Luke, J.J. & Gajewski, T.F. (2018). The commensal microbiome is associated with anti-PD-1 efficacy in metastatic melanoma patients. *Science*, 359(6371), 104–108. DOI: 10.1126/science.aao3290
- [REF:routy2018] Routy, B., Le Chatelier, E., Derosa, L., Duong, C.P.M., Alou, M.T., Daillère, R., Fluckiger, A., Messaoudene, M., Rauber, C., Roberti, M.P., Fidelle, M., Flament, C., Poirier-Colame, V., Ber, P.E., Tber, C., Sber, A.R., Ibrahim, Y.S.K., Eber, J., Goubert, R., Zitvogel, L. et al. (2018). Gut microbiome influences efficacy of PD-1-based immunotherapy against epithelial tumors. *Science*, 359(6371), 91–97. DOI: 10.1126/science.aan3706
- [REF:mager2021] Mager, L.F., Burkhard, R., Pett, N., Cooke, N.C.A., Brown, K., Ramber, H., Mucber, B., Wander, S.E., Kber, M., Bber, J., MacPherson, A.J., McCoy, K.D. & Hapfelmeier, S. (2021). Microbiome-derived inosine modulates response to checkpoint inhibitor immunotherapy. *Science*, 369(6510), 1481–1489. DOI: 10.1126/science.abc3421
- [REF:baruch2021] Baruch, E.N., Youngster, I., Ben-Betzalel, G., Ortenberg, R., Lahat, A., Katz, L., Adler, K., Dick-Necula, D., Raskin, S., Bloch, N., Rober, D., Shapira-Frommer, R., Brber, M., Mber, J., Ben-Aharon, I., Ziv, O., Khan, M.A.W., Amit, M., Cooks, T., Berber, S., Schachter, J., Merbl, Y., Geva-Zatorsky, N. & Markel, G. (2021). Fecal microbiota transplant promotes response in immunotherapy-refractory melanoma patients. *Science*, 371(6529), 602–609. DOI: 10.1126/science.abb5920
- [REF:davar2021] Davar, D., Dzuber, A.K., Tber, H., Sber, J.C., Brber, B.F., Anber, S., Eber, R.A.N., Lber, A., Luke, J.J., Divakar, P., Bui, K., Kalac, M., Ber, A.P., Mber, M.E., Hernandez, A.V., Zber, M.A., Tumber, T.C., Nber, C.E., Mber, M.J., Gber, J.M., Tber, E.J., Kirkwood, J.M., Baldo, C., Gharaibeh, R.Z., Zarour, H.M. & Jobin, C. (2021). Fecal microbiota transplant overcomes resistance to anti-PD-1 therapy in melanoma patients. *Science*, 371(6529), 595–602. DOI: 10.1126/science.abf3363

Additional references to be verified and expanded during manuscript preparation.

Chapter 13: The Gut-Brain Axis — Microbes and the Mind

You have probably heard, at some point in your life, the advice to “trust your gut.” Perhaps before a job interview, a first date, or a difficult decision, someone told you that your intestines somehow knew things your conscious mind did not. The phrase is usually offered as folk wisdom — a vague nod to intuition, not meant to be taken literally.

It turns out to be more literal than anyone suspected.

Over the past fifteen years, a rapidly growing body of research has revealed that the gut and the brain are in constant, bidirectional communication — exchanging signals through nerves, hormones, immune molecules, and the chemical products of microbial metabolism. This conversation is not a metaphor. It is a measurable, manipulable, physiologically consequential dialogue conducted through defined molecular pathways. And the microbiome — the trillions of bacteria, fungi, and viruses that inhabit the gastrointestinal tract — is not merely a passive bystander in this dialogue. It is an active participant, producing molecules that influence how we respond to stress, how we feel, and possibly even how we think.

In the previous chapter, we saw how the microbiome shapes the immune system — training it, calibrating it, negotiating the terms of coexistence through a molecular conversation refined over hundreds of millions of years. Now we follow a different communication channel: the one that runs between the gut and the brain. If the immune negotiation is about distinguishing friend from foe, the gut-brain conversation is about something more intimate — mood, anxiety, cognition, and the murky territory where the body’s internal state becomes what we experience as emotion.

This is also, it must be said at the outset, one of the most overhyped areas of microbiome science. The popular press has seized on early findings with an enthusiasm that far outstrips the evidence, producing headlines that suggest your gut bacteria control your personality, cause your depression, and can be fixed with the right yoghurt. The reality is more nuanced, more interesting, and more honest. Some of the findings in this field are among the most provocative in modern biology. Others are preliminary, inconsistent, or have failed to replicate in larger studies. We will try to distinguish between them as we go — because the gut-brain axis deserves better than hype. It deserves careful attention.

13.1 The Vagus Nerve as an Information Superhighway

The Wanderer

To understand how the gut talks to the brain, we need to start with a nerve — the longest and most wide-ranging nerve in the human body. The **vagus nerve** (from the Latin *vagus*, meaning “wandering”) is the tenth cranial nerve, and it deserves its name. It originates in the brainstem, in a cluster of neurons called the **dorsal motor nucleus** and the **nucleus tractus solitarius**, and from there it wanders south — branching through the neck, the chest, and the abdomen, sending fibres to the heart, the lungs, the oesophagus, the stomach, the liver, the pancreas, and the entire length of the intestine down to the proximal two-thirds of the colon.

Think of it as a two-lane highway connecting the brain to the body's major organs. About eighty per cent of the fibres in the vagus nerve are **afferent** — meaning they carry information *upward*, from the organs to the brain. The remaining twenty per cent are **efferent** — carrying instructions *downward*, from the brain to the organs. This ratio is itself telling. The vagus nerve is not primarily a command cable through which the brain issues orders to the gut. It is, overwhelmingly, a sensory channel — a listening device through which the brain monitors what is happening below.

And what is happening below is enormously complex. The gut contains its own nervous system — the **enteric nervous system (ENS)**, which we first encountered in Chapter 5. The ENS consists of somewhere between 200 and 600 million neurons embedded in the walls of the gastrointestinal tract, organised into two major nerve networks (the myenteric plexus and the submucosal plexus) that regulate gut motility, secretion, and blood flow largely independently of the brain. The ENS is sometimes called the “second brain,” and while the label is an oversimplification — the ENS cannot compose poetry or do arithmetic — it captures an important truth: the gut has enough neural processing power to manage most of its own affairs without consulting the brain at all. The vagus nerve is the trunk line that connects these two nervous systems, allowing each to influence the other.

For gut-brain communication, the critical players are the afferent vagal fibres that terminate in the gut wall. These nerve endings sit in the mucosa and submucosa of the intestine — close to the epithelial lining, close to the immune cells of the lamina propria, and, crucially, close to the microbial ecosystem on the other side of the epithelial barrier. The vagal afferents do not make direct contact with gut bacteria. But they are exquisitely sensitive to the chemical environment around them. When gut bacteria produce metabolites — short-chain fatty acids, neurotransmitters, tryptophan derivatives — these molecules can diffuse across the epithelium and activate receptors on vagal nerve endings, sending electrical signals up to the brainstem. From there, the signals are relayed to higher brain regions involved in mood, emotion, and cognition — including the amygdala (the brain's threat-detection centre), the hippocampus (memory), and the prefrontal cortex (decision-making and emotional regulation).

The vagus nerve, in other words, gives the gut microbiome a direct line to the emotional brain.

The Experiment That Proved It

The importance of vagal signalling in the microbiome-brain connection was demonstrated with remarkable clarity in a 2011 study by Javier Bravo, John Cryan, and colleagues at University College Cork — a group that has done more than perhaps any other to define the field of microbiome-brain research [REF:bravo2011].

Bravo's team fed mice a strain of *Lactobacillus rhamnosus* (JB-1) — a common bacterium found in fermented dairy products — and observed a striking suite of changes. The treated mice showed reduced anxiety-like behaviour in standard laboratory tests: they were more willing to explore open, exposed areas of a maze (which mice normally avoid, because open spaces mean vulnerability to predators) and spent more time in illuminated areas of a light-dark box. They also showed reduced depressive-like behaviour in a forced swim test — a widely used (if imperfect) assay in which the amount of time a mouse spends immobile in a beaker of water is taken as a proxy for behavioural despair. The bacterial treatment lowered levels of **corticosterone** — the rodent equivalent of the human stress hormone cortisol — in the blood after a stressful experience.

At the molecular level, the changes were equally striking. The *L. rhamnosus*-fed mice showed altered expression of **GABA receptors** across multiple brain regions — a finding with direct rel-

evance to anxiety and depression, since GABA (gamma-aminobutyric acid) is the brain's principal inhibitory neurotransmitter, and disruptions in GABA signalling are implicated in both conditions. The pattern of changes — increased GABA-B receptors in some cortical regions, decreased levels in the hippocampus and amygdala — was region-specific and consistent with reduced anxiety.

Then came the critical control. Bravo's team repeated the entire experiment in mice that had undergone a **vagotomy** — a surgical procedure in which the vagus nerve is cut. In the vagotomised mice, every single effect of the *Lactobacillus* treatment disappeared. The bacteria were still in the gut. They were still alive, still metabolising, still producing their chemical products. But without the vagus nerve to carry the signal upward, none of those products reached the brain. The anxiety reduction vanished. The corticosterone changes vanished. The GABA receptor alterations vanished. The vagus nerve was not merely one of several communication channels — for this particular bacterium, it was the *only* channel that mattered.

This was a watershed result. It established, with a clean surgical control, that the vagus nerve is the essential conduit for at least some forms of microbiome-to-brain communication. And it raised a provocative possibility: if a single bacterial species, delivered orally, could rewire GABA receptor expression across the brain via the vagus nerve, what was the entire gut microbiome doing to the brain, all the time, through the same pathway?

Beyond the Vagus: Other Communication Routes

The vagus nerve is the most direct and best-studied route from gut to brain, but it is not the only one. The microbiome communicates with the central nervous system through at least four major channels, which operate simultaneously and often interact:

The neural pathway — the vagus nerve and the enteric nervous system, as just described. Also includes the spinal afferent nerves, which carry signals from the gut to the spinal cord and from there to the brain, particularly pain and discomfort signals.

The immune pathway — which we covered in the previous chapter. Gut bacteria and their products activate immune cells in the gut wall, which release cytokines (chemical messengers of the immune system) into the bloodstream. Some of these cytokines — particularly pro-inflammatory ones like interleukin-1 β (IL-1 β), interleukin-6 (IL-6), and tumour necrosis factor alpha (TNF- α) — can cross the blood-brain barrier or signal through it, influencing brain function and behaviour. This is why you feel listless, irritable, and mentally foggy when you have an infection: the “sickness behaviour” that accompanies fever is not a direct effect of the pathogen on the brain but a consequence of inflammatory cytokines acting on brain circuits that regulate mood and motivation.

The endocrine pathway — through the production of hormones and hormone-like molecules. The gut is the body's largest endocrine organ, producing more than twenty distinct hormones from specialised cells scattered throughout the intestinal lining. The microbiome influences many of these hormones — and it also produces its own neuroactive compounds, as we are about to see.

The metabolic pathway — through the production of small molecules that enter the bloodstream and reach the brain. Short-chain fatty acids (SCFAs), which we discussed in Chapters 5 and 12, are one example: they are produced by bacterial fermentation of dietary fibre, absorbed into the bloodstream, and can cross the blood-brain barrier. Tryptophan metabolites — produced when gut bacteria break down the amino acid tryptophan — are another, and they activate receptors in the brain that influence mood and cognition. We will meet these metabolites again shortly.

These four pathways — neural, immune, endocrine, and metabolic — constitute the **gut-brain axis**, a term that has become one of the most-used phrases in microbiome science. When the microbial dimension is explicitly included, some researchers prefer the term **microbiota-gut-brain axis** — an ungainly phrase, but one that captures the essential point: the communication is not merely between the gut and the brain but involves the microbial ecosystem as a third party in the conversation.

13.2 Microbial Production of Neurotransmitters: Serotonin, GABA, Dopamine

The Chemical Factory in Your Gut

Here is a fact that surprises almost everyone who encounters it for the first time: roughly ninety per cent of the serotonin in your body is not in your brain. It is in your gut.

Serotonin — or, to give it its chemical name, 5-hydroxytryptamine (5-HT) — is one of the most famous molecules in neuroscience. It is the target of the selective serotonin reuptake inhibitors (SSRIs) that are the first-line treatment for depression. It is involved in regulating mood, sleep, appetite, and cognition. It is, in the popular imagination, the “happiness chemical.” And the vast majority of it is produced not by neurons in the brain but by a specialised type of cell in the lining of the intestine called the **enterochromaffin cell** — named for its affinity for chromium salts, which stain it a distinctive brown under the microscope.

Enterochromaffin cells are scattered throughout the gut epithelium, and they function as chemical sensors. They detect signals from the gut lumen — including signals from bacteria — and respond by releasing serotonin into the surrounding tissue. From there, the serotonin acts on nearby nerve endings (including vagal afferents), on immune cells, on smooth muscle cells that control gut motility, and on platelets in the bloodstream, which absorb it and carry it throughout the body.

In 2015, Elaine Hsiao and colleagues at the California Institute of Technology published a landmark study in *Cell* that quantified the microbiome’s role in this process [REF:yano2015]. Working with germ-free mice, they showed that the gut microbiota was responsible for approximately sixty per cent of the serotonin produced in the colon. Germ-free mice had dramatically lower levels of both colonic and circulating serotonin compared with conventionally raised mice. When germ-free mice were colonised with a consortium of about twenty species of spore-forming bacteria — predominantly Clostridia — serotonin levels were restored to near-normal. The bacteria were not producing serotonin themselves (at least not in significant quantities). Instead, they were *signalling* to the enterochromaffin cells to ramp up production, by releasing metabolites — including short-chain fatty acids and secondary bile acids — that activated serotonin-synthesis pathways in the host cells.

The implications were profound but also easily overstated, so let us be precise about what this study did and did not show. It demonstrated that the gut microbiome has a powerful influence on peripheral serotonin — the serotonin circulating in the blood and acting in the gut. It did *not* demonstrate that the microbiome directly controls serotonin in the brain. And this distinction matters enormously, because peripheral serotonin and brain serotonin are, for practical purposes, two separate pools.

The Blood-Brain Barrier Problem

Here we encounter a critical piece of biology that is often glossed over in popular accounts of the gut-brain axis. Serotonin does not cross the blood-brain barrier. The serotonin in your gut and the serotonin in your brain are produced independently, by different cells, and they cannot exchange between the two compartments. The same is true for most other neurotransmitters: GABA, dopamine, and norepinephrine are all produced by gut bacteria or by gut cells under microbial influence, but none of them can simply diffuse from the intestine to the brain.

This is not a minor detail. It is the central puzzle of gut-brain communication. If the gut is producing enormous quantities of neurotransmitters that cannot reach the brain directly, how does the microbiome influence brain function?

The answer — and this is where the picture becomes genuinely interesting rather than simplistically reductive — is that the microbiome influences the brain through *indirect* pathways, each of which is more subtle and more informative than a simple chemical pipeline would be.

Pathway one: vagal signalling. Serotonin released by enterochromaffin cells does not need to travel to the brain to influence it. It acts *locally*, on vagal afferent nerve endings that are right there in the gut wall. Recent research has clarified that this communication happens not through conventional synaptic transmission but through **diffusion** — the serotonin seeps through the tissue until it reaches nerve endings in the lamina propria, where it activates serotonin receptors (particularly the 5-HT₃ receptor) on the vagal fibres. The vagal fibres then transmit an electrical signal to the brainstem — specifically to the nucleus tractus solitarius — which relays the information to the dorsal raphe nucleus (where the brain's own serotonergic neurons reside) and to the locus coeruleus (which produces norepinephrine). The gut's serotonin, in this model, is not a brain neurotransmitter. It is a *signal* — a message that travels by nerve rather than by blood, using the vagus as its telephone line.

Pathway two: the precursor route. While serotonin cannot cross the blood-brain barrier, its precursor can. Serotonin is synthesised from the amino acid **tryptophan**, which is obtained from dietary protein and *does* cross the blood-brain barrier via a specific amino acid transporter. Once inside the brain, tryptophan is converted to serotonin by local neurons. The gut microbiome influences this pathway by competing for tryptophan. Certain gut bacteria metabolise tryptophan through alternative pathways — converting it into indole, indole-3-propionic acid, and other derivatives rather than allowing it to be absorbed and sent to the brain. Other bacteria promote tryptophan absorption. The net effect is that the composition of the gut microbiome can alter how much tryptophan reaches the brain, and therefore how much serotonin the brain can produce. Think of it as a competition for raw materials: the gut microbiome and the brain are both drawing from the same pool of tryptophan, and the microbiome gets first access.

Pathway three: short-chain fatty acids. Unlike serotonin, the SCFAs produced by bacterial fermentation of dietary fibre — acetate, propionate, and butyrate — *can* cross the blood-brain barrier, albeit in modest quantities. Once in the brain, butyrate in particular has been shown to influence gene expression by inhibiting **histone deacetylases** (enzymes that regulate how tightly DNA is wound around its protein scaffolding, and therefore which genes are active). Butyrate also increases the production of **brain-derived neurotrophic factor (BDNF)** — a protein critical for neuronal survival, growth, and plasticity — in the hippocampus. Reduced BDNF is one of the most consistent findings in depression research, and the fact that a microbial metabolite can influence its levels in the brain is one of the more tantalising links in the gut-brain story.

Pathway four: immune signalling. As we discussed in Chapter 12, the microbiome continuously modulates the immune system, and immune signals reach the brain. When the gut microbiome is disrupted — by antibiotics, poor diet, or stress — the resulting low-grade inflammation produces cytokines that act on the brain, altering mood, motivation, and cognition. This is not gut-brain communication in the romantic sense of the term. It is closer to the brain receiving a status report about trouble in the periphery and adjusting its operations accordingly — dampening motivation to conserve energy, increasing vigilance, and reducing social behaviour. The overlap between the symptoms of chronic low-grade inflammation and the symptoms of depression (fatigue, anhedonia, social withdrawal, difficulty concentrating) is not coincidental. It may reflect a shared biological mechanism.

GABA: The Calming Signal

If serotonin is the most famous neurotransmitter in the gut-brain story, GABA (gamma-aminobutyric acid) is arguably the most relevant to anxiety. GABA is the brain's principal inhibitory neurotransmitter — the chemical signal that tells neurons to slow down, stop firing, and calm the circuit. Every major class of anti-anxiety medication — the benzodiazepines (diazepam, lorazepam), the barbiturates, even alcohol — works by enhancing GABA signalling in the brain. When GABA circuits are underactive, the result is anxiety, hypervigilance, and difficulty relaxing. When they are overactive, the result is sedation.

In 2019, Philip Strandwitz and colleagues at Northeastern University published an elegant study in *Nature Microbiology* demonstrating that GABA production is a common feature of the gut microbiome [REF:strandwitz2019]. The study began with a puzzle: the team had isolated a gut bacterium, designated KLE1738, that refused to grow in any standard culture medium. After extensive troubleshooting, they discovered that KLE1738 required GABA as a growth factor — it could not survive without it. When they co-cultured KLE1738 with *Bacteroides fragilis* (the same species we met in Chapter 12 for its immune-modulatory polysaccharide A), growth was restored. The reason: *Bacteroides fragilis* produces GABA in substantial quantities, using an enzyme called glutamate decarboxylase to convert the amino acid glutamate into GABA.

The team went on to show that GABA production is widespread among gut bacteria, with *Bacteroides* species being particularly prolific producers. They then examined metagenomic data from human stool samples and found that the relative abundance of GABA-producing bacteria in the gut correlated with brain signatures relevant to depression — a finding that was consistent with the larger population studies we will discuss in section 13.5.

Like serotonin, GABA produced in the gut does not cross the blood-brain barrier directly. But it does not need to. GABA receptors are expressed on intestinal epithelial cells and on vagal afferent nerve endings in the gut wall. When gut bacteria produce GABA, it acts locally — binding to these receptors, modulating intestinal function, and triggering nerve signals that travel up the vagus to the brain. This is precisely the mechanism demonstrated in the Bravo *L. rhamnosus* study described in section 13.1: the bacterium's effects on brain GABA receptor expression depended entirely on an intact vagus nerve. The gut bacterium was not sending GABA to the brain. It was sending a *message* about GABA, via the nerve.

Dopamine, Norepinephrine, and the Broader Chemical Repertoire

The microbial chemical factory does not stop at serotonin and GABA. Gut bacteria have been shown to produce or modulate levels of a remarkably wide range of neuroactive compounds:

Dopamine — the neurotransmitter associated with reward, motivation, and pleasure — is produced by several gut bacterial species, including *Lactobacillus* and *Enterococcus*. The enzyme responsible is typically **tyrosine decarboxylase**, which converts the amino acid tyrosine into tyramine and, under some conditions, can convert L-DOPA (the precursor to dopamine) into dopamine itself. As with serotonin, gut-produced dopamine acts primarily in the periphery — regulating gut motility and intestinal immune function — rather than reaching the brain directly.

Norepinephrine (also called noradrenaline) — the neurotransmitter of alertness and the fight-or-flight response — is produced by certain *Escherichia*, *Bacillus*, and *Saccharomyces* species. Its gut-level effects include modulating intestinal motility and influencing the growth of certain bacterial species, suggesting a feedback loop in which the host's stress response (which releases norepinephrine into the gut) alters the microbiome, and the microbiome's own norepinephrine production further modifies the local environment.

Acetylcholine — the neurotransmitter of the parasympathetic nervous system, involved in rest, digestion, and memory — is produced by *Lactobacillus* species, among others.

Histamine — best known for its role in allergic reactions but also a neurotransmitter involved in wakefulness and appetite — is produced by several gut bacteria, including *Lactobacillus reuteri* and certain strains of *Morganella* and *Klebsiella*.

The sheer breadth of this chemical repertoire is astonishing. The gut microbiome produces, at last count, virtually every major class of neurotransmitter found in the human brain. This does not mean the microbiome is “controlling” the brain — the vast majority of these molecules act locally, within the gut, and their influence on brain function is indirect, mediated through the neural, immune, and endocrine pathways described above. But it does mean that the gut is not merely a digestive organ. It is a chemical factory operating at industrial scale, producing molecules that, through various intermediaries, can influence how the brain works.

The question — and it is a question that the field has not yet fully answered — is how much this matters. Does the microbiome's neurotransmitter production have a measurable, clinically significant effect on human mood, cognition, and mental health? Or is it a biological curiosity — real but trivial in the context of the brain's own vastly more powerful neurotransmitter systems? The next four sections will examine the evidence.

13.3 The HPA Axis, Stress, and Microbial Modulation of Cortisol

The Stress Thermostat

When you encounter something stressful — a near-miss in traffic, an argument with a colleague, a sudden loud noise — your body launches a coordinated physiological response that prepares you to fight or flee. Your heart rate increases. Your pupils dilate. Your muscles tense. Glucose floods into the bloodstream. And all of this is orchestrated by a hormonal cascade called the **hypothalamic-pituitary-adrenal axis**, or HPA axis — the body's central stress-response system.

The HPA axis works like a three-stage amplifier. First, the **hypothalamus** — a small structure at the base of the brain that serves as the body's master control centre for hormones — detects the stressor and releases a signalling molecule called **corticotropin-releasing hormone (CRH)**. CRH travels a short distance through a dedicated blood vessel to the **anterior pituitary gland**,

which sits just below the hypothalamus. The pituitary responds by releasing **adrenocorticotrophic hormone (ACTH)** into the general bloodstream. ACTH travels to the **adrenal glands**, which sit like small caps atop the kidneys, and stimulates them to produce **cortisol** — the body's primary stress hormone (in rodents, the equivalent is corticosterone).

Cortisol does many things. It mobilises energy reserves, suppresses non-essential functions (including digestion, immune activity, and reproduction), heightens alertness, and prepares the body for action. It also, critically, feeds back to the hypothalamus and pituitary to shut down its own production — a negative feedback loop that ensures the stress response does not run indefinitely. Think of it as a thermostat: the hypothalamus sets the desired temperature, cortisol is the heat output, and when the room gets warm enough, the system turns itself down. A well-functioning HPA axis, like a well-functioning thermostat, hits the target and stops. A dysregulated one keeps pumping out heat long after the room is sweltering.

This is a beautifully designed system for dealing with acute threats — a predator, a natural disaster, a physical confrontation. It becomes a problem when the stress is chronic. When the stressor does not go away — when it is a difficult job, a troubled relationship, financial insecurity, or the grinding baseline anxiety of modern life — the HPA axis remains activated, cortisol levels stay elevated, and the negative feedback loop begins to lose its effectiveness. Chronic cortisol elevation damages the hippocampus (impairing memory), suppresses the immune system (increasing susceptibility to infection), disrupts gut barrier function (increasing intestinal permeability), and alters the composition of the gut microbiome.

And here the story becomes circular, because the gut microbiome, it turns out, helps calibrate the HPA axis in the first place.

The Germ-Free Mice That Couldn't Handle Stress

The foundational study linking the gut microbiome to stress came in 2004, from Nobuyuki Sudo and colleagues at Kyushu University in Japan — a study published before the term “gut-brain axis” had entered common scientific parlance [REF:sudo2004].

Sudo's team subjected germ-free mice — animals raised without any microbial exposure — to a standard laboratory stressor: restraint stress, in which the mouse is placed in a small tube that prevents movement for a defined period. They then measured the animals' hormonal response and compared it with that of conventionally raised mice exposed to the same stressor.

The results were dramatic. The germ-free mice showed a massively exaggerated stress response — their levels of ACTH and corticosterone shot up far higher than those of conventional mice, and took much longer to return to baseline. Their HPA axis, without microbial input, was stuck on high. The stress thermostat was broken — or, more precisely, it had never been properly calibrated.

Sudo then asked whether this could be corrected. He colonised germ-free mice with a single species: *Bifidobacterium infantis* — a bacterium that, as we saw in Chapter 3, is one of the earliest colonisers of the infant gut, nourished by the oligosaccharides in breastmilk. The *Bifidobacterium* treatment normalised the stress response, bringing ACTH and corticosterone levels back to the range seen in conventionally raised animals.

But there was a catch — and it was an important one. The normalisation only worked if the bacteria were introduced early in life. When Sudo colonised germ-free mice at six weeks of age (roughly equivalent to early adolescence in human terms), the stress response was rescued. When he waited

until fourteen weeks (young adulthood), it was not. There was a **critical developmental window** during which microbial colonisation had to occur for the HPA axis to develop normally. Miss the window, and the damage was, at least in this experimental system, permanent.

This finding echoed what we learned in Chapter 12 about the weaning reaction and immune development: there are specific periods in early life during which the microbiome must be present for normal physiological development. The immune system has its window. The stress response has its window. And the microbiome is the clock that sets them both.

The Vicious Cycle

The relationship between stress and the microbiome is not one-directional. Stress does not merely *result from* a disrupted microbiome. It also *causes* microbiome disruption — creating a feedback loop that can become self-reinforcing.

When the HPA axis is activated and cortisol levels rise, several things happen in the gut. The intestinal barrier becomes more permeable — the tight junctions between epithelial cells loosen, a phenomenon sometimes described, with deliberate oversimplification, as “leaky gut.” The mucus layer thins. The composition of the microbial community shifts, typically toward a less diverse state with a reduction in beneficial species (particularly SCFA-producing Clostridia and *Bifidobacterium*) and an expansion of potentially inflammatory species (particularly certain Proteobacteria). The immune tone of the gut changes, with increased pro-inflammatory cytokine production.

These changes, in turn, feed back to the brain. A more permeable gut allows bacterial products — particularly **lipopolysaccharide (LPS)**, the inflammatory outer-membrane component of Gram-negative bacteria that we discussed in Chapter 12 — to leak into the bloodstream. Even small amounts of circulating LPS activate the innate immune system, producing a low-grade systemic inflammation that affects the brain. Circulating pro-inflammatory cytokines cross the blood-brain barrier (or signal through it) and activate microglia — the brain’s resident immune cells — producing neuroinflammation that further dysregulates the HPA axis.

The result is a vicious cycle — a biological version of the audio feedback loop you hear when a microphone is placed too close to its own speaker. Stress disrupts the microbiome □ the disrupted microbiome produces inflammatory signals □ the inflammatory signals reach the brain □ the brain mounts a further stress response □ the stress response disrupts the microbiome further. Each turn of the cycle amplifies the signal. Breaking this loop — whether through stress reduction, dietary intervention, or microbiome restoration — is one of the therapeutic goals that motivates much of the current research on the gut-brain axis.

13.4 Animal Studies: Anxious Mice, Bold Mice, and Faecal Transplants

Personality in a Pellet

If the Sudo study showed that the microbiome calibrates the stress response, a study published seven years later showed something even more startling: the microbiome can influence personality.

In 2011, Premysl Bercik, Stephen Collins, and colleagues at McMaster University in Canada performed an experiment that has become one of the most cited — and most discussed — in the entire gut-brain field [REF:bercik2011].

The experiment exploited a well-known difference between two inbred mouse strains. **BALB/c mice** are naturally timid. They avoid open spaces, stick to the walls of enclosures, and display high levels of anxiety-like behaviour in standard laboratory tests. **NIH Swiss mice**, by contrast, are bold — adventurous, exploratory, and relatively unfazed by novel environments. These behavioural differences are stable, reproducible, and have traditionally been attributed to genetic variation between the strains. BALB/c mice are anxious because their genes make them anxious. NIH Swiss mice are bold because their genes make them bold. That was the assumption.

Bercik's team tested that assumption by swapping the microbiomes. They raised both strains as germ-free animals — born and kept in sterile isolators, with no microbial colonisation. Then they colonised germ-free BALB/c mice with the gut microbiota from NIH Swiss mice, and germ-free NIH Swiss mice with the gut microbiota from BALB/c mice.

The results were remarkable. The BALB/c mice — the genetically “anxious” strain — colonised with bold-mouse microbiota became *bolder*. They explored more, avoided less, and showed behavioural patterns more reminiscent of NIH Swiss mice than of their own genetic background. Conversely, the NIH Swiss mice — the genetically “bold” strain — colonised with anxious-mouse microbiota became *more cautious*. Their exploratory behaviour decreased. Their anxiety-like behaviour increased. The behavioural phenotype — the personality, if you will — had transferred with the microbiome.

At the molecular level, the changes correlated with alterations in **brain-derived neurotrophic factor (BDNF)** — a protein that supports neuronal growth, survival, and the formation of new synaptic connections. BDNF levels in the hippocampus tracked the behavioural changes: mice that became bolder had higher hippocampal BDNF, and mice that became more anxious had lower levels. BDNF is one of the most consistently disrupted molecules in depression and anxiety disorders, and the finding that the microbiome could influence its levels in a specific brain region added mechanistic weight to the behavioural observations.

To be clear about what this study showed and what it did not: the microbiome did not completely override genetics. The colonised BALB/c mice did not become indistinguishable from wild-type NIH Swiss mice. The genetic background still mattered. But the microbiome pushed the animals' behaviour along a spectrum — shifting the dial in a direction that was consistent and reproducible. The personality was not entirely in the genes. Part of it was in the gut.

The Germ-Free Behavioural Profile

The Bercik study was part of a broader body of work, emerging around the same period, that characterised the behavioural profile of germ-free mice — and the picture that emerged was consistently abnormal.

In 2011, Rochellys Diaz Heijtz and colleagues at the Karolinska Institute in Stockholm published a study in *PNAS* showing that germ-free mice displayed increased motor activity and reduced anxiety-like behaviour compared with conventionally raised controls [REF:diazheijt2011]. The germ-free mice were more active, more willing to explore open and exposed areas, and less cautious in novel environments. At the molecular level, they showed altered expression of genes involved in synaptic plasticity, neurotransmitter signalling, and second-messenger pathways in the striatum — a brain region involved in motor control and reward-seeking behaviour.

This might sound paradoxical. If the microbiome calms anxiety (as the Bravo *L. rhamnosus* study suggested), why would germ-free mice — which lack a microbiome entirely — be *less* anxious? The

answer lies in the distinction between a properly calibrated system and no system at all. The microbiome does not simply reduce anxiety. It calibrates the anxiety response — setting it to an appropriate level for the animal’s environment. A germ-free mouse is not “calm.” It is *under-responsive* — lacking the normal caution that would keep a wild mouse alive. Its reduced anxiety is not a sign of wellbeing but of a brain that has not been properly tuned.

Germ-free mice also show profound deficits in social behaviour. They spend less time interacting with unfamiliar mice, show reduced preference for social novelty (normally, mice are intensely curious about strangers), and display abnormal patterns of social investigation. These social deficits have been linked to altered signalling in brain regions associated with social cognition, including the prefrontal cortex and the amygdala — and, intriguingly, to changes in the neuropeptide **oxytocin**, which is sometimes called the “social bonding hormone” for its role in attachment, trust, and social recognition.

Recent work from Sarkis Mazmanian’s group at Caltech has begun to identify the specific microbial metabolites responsible for these social effects. In studies published in 2024 and 2025, the team showed that **hippuric acid** — a metabolite produced by gut bacterial processing of dietary polyphenols — can rescue social behaviour in germ-free mice when administered directly. Hippuric acid appears to act by increasing oxytocin expression in the hypothalamic paraventricular nucleus — a finding that provides a remarkably specific molecular link between a microbial product and a social behaviour, mediated by a well-characterised neuropeptide system.

What the Animals Tell Us — And What They Don’t

The animal studies on the microbiome-brain connection are, collectively, some of the most striking results in modern neuroscience. A single bacterial species can rewire GABA receptor expression across the brain. A microbiome transplant can shift an animal’s behavioural profile along the anxiety spectrum. Germ-free mice are socially impaired, stress-hyper-reactive, and behaviourally abnormal — and all of these deficits can be at least partially corrected by microbial colonisation.

But we must handle these findings with care, for three reasons.

First, **mice are not people**. The mouse brain and the human brain differ enormously in complexity, size, and the relative importance of cortical processing (which is far greater in humans). Behaviours that serve as proxies for anxiety or depression in mice — time spent immobile in a forced swim test, willingness to eat in an unfamiliar environment — are, at best, crude analogues of the human experiences they are meant to model. A mouse that floats motionless in a beaker of water may be “behaviourally despairing” in some sense, but it is also conserving energy, and the relationship between that behaviour and the lived experience of human depression is tenuous.

Second, **germ-free animals are profoundly artificial**. A mouse raised in complete sterility from birth has not merely lost its microbiome. It has developed in a condition that has never existed in the natural history of the species. Its immune system is underdeveloped (as we saw in Chapter 12). Its gut barrier is abnormal. Its metabolism is altered. The behavioural changes in germ-free mice may reflect the absence of the microbiome specifically, or they may reflect the cascading consequences of developing without an immune system, without normal intestinal function, and without any of the metabolites that the microbiome provides. Separating these effects is technically challenging and, in many studies, has not been attempted.

Third, **translation to humans is uncertain**. The gap between showing that a bacterial metabolite changes mouse behaviour and showing that it has a meaningful effect on human mental health

is vast. Humans have far more complex social environments, far more variable diets, far more diverse microbiomes, and — most importantly — subjective experiences that cannot be measured by putting them in a maze. The animal studies provide *proof of concept* that the microbiome can influence brain function through defined pathways. They do not tell us how much it matters in the context of human psychiatric illness.

With those caveats firmly in mind, let us turn to what we know about the microbiome-brain connection in humans.

13.5 Depression, Anxiety, and the Emerging Concept of “Psychobiotics”

A New Word for an Old Idea

In 2013, Timothy Dinan and John Cryan — the same researchers whose work with *Lactobacillus rhamnosus* and the vagus nerve helped launch the gut-brain field — published a paper in *Biological Psychiatry* that introduced a term destined to escape from the scientific literature into the popular imagination: **psychobiotics** [REF: dinan2013].

Dinan and Cryan defined a psychobiotic as “a live organism that, when ingested in adequate amounts, produces a health benefit in patients suffering from psychiatric illness.” The term was deliberately modelled on the word “probiotic,” and it carried the same mix of scientific ambition and marketing appeal. The idea was simple and seductive: if the gut microbiome influences the brain, and if specific bacteria can alter mood and anxiety in animal models, then perhaps carefully selected bacterial strains — psychobiotics — could be developed as treatments for depression, anxiety, and other mental health conditions. Not replacing conventional therapy, but supplementing it. A pill containing bacteria rather than serotonin reuptake inhibitors.

The concept was not without foundation. By 2013, the animal evidence was substantial: *L. rhamnosus* reduced anxiety via the vagus nerve; *Bifidobacterium infantis* normalised the stress response in germ-free mice; *Bifidobacterium longum* reduced anxiety-like behaviour in a mouse model of colitis; and various *Lactobacillus* and *Bifidobacterium* strains had shown effects on GABA, serotonin, and BDNF pathways. The question was whether any of this translated to humans — whether bacteria ingested by a person with depression or anxiety would produce meaningful clinical improvement.

The honest answer, a decade later, is: probably, but modestly, and we are still working out the details.

From Mice to Humans: The Transfer Experiment

The most direct bridge between the animal models and human depression came in 2016, when John Kelly, Cryan, Dinan, and colleagues published a study that applied the microbiome-transfer paradigm — so successful in mice — to human depression [REF: kelly2016].

The team collected faecal samples from 34 patients with major depressive disorder and 33 healthy controls, and transplanted the microbiota into rats whose own gut bacteria had been depleted by antibiotics. The rats receiving microbiota from depressed donors developed behavioural changes consistent with a depressive-like state: reduced preference for a sucrose solution (a standard measure of anhedonia — the inability to experience pleasure, which is one of the core symptoms of de-

pression), increased anxiety-like behaviour, and altered tryptophan metabolism. The rats receiving microbiota from healthy donors showed no such changes.

This was an important study because it demonstrated — in a mammalian system, using human donor material — that the microbial community associated with depression could, by itself, induce depressive-like changes in a recipient animal. It moved the conversation from correlation (“depressed people have different gut bacteria”) to something closer to causation (“the gut bacteria of depressed people can produce depression-relevant changes in another organism”). But it was still an animal study, with all the limitations that implies.

The Flemish Gut Flora Project: Population-Scale Evidence

The largest and most carefully controlled study linking specific gut bacteria to human mental health was published in 2019 by Mireia Valles-Colomer and colleagues, using data from the **Flemish Gut Flora Project** — a population-based cohort of over 1,000 individuals in Belgium with detailed health, lifestyle, and microbiome data [REF:vallescolomer2019].

Valles-Colomer’s team took a systematic approach. Rather than starting with a hypothesis about which bacteria might matter, they examined the entire gut microbiome in relation to participants’ quality of life and depression diagnoses (using both self-reported measures and GP diagnoses), while carefully controlling for confounders including age, sex, BMI, stool consistency, and — crucially — antidepressant use, which can itself alter the microbiome.

Two findings stood out. First, the butyrate-producing genera **Faecalibacterium** and **Coprococcus** were consistently associated with higher quality-of-life indicators. Second, **Coprococcus** and **Dialister** were significantly depleted in participants with depression — a finding that held even after excluding participants taking antidepressants, ruling out the possibility that the medication was driving the microbial differences rather than the depression itself.

The team then validated their findings in an independent cohort — the LifeLines DEEP study in the Netherlands, comprising over 1,000 additional participants — and found the same associations. Coprococcus and Dialister depletion in depression replicated across two countries, two research groups, and two independent analytical pipelines.

Valles-Colomer’s group also performed something novel: they mined the genomes of the gut bacteria for their capacity to produce neuroactive compounds. They found that bacterial pathways for producing the dopamine metabolite **3,4-dihydroxyphenylacetic acid (DOPAC)** correlated positively with mental quality of life, and that GABA synthesis pathways were inversely associated with depression. The bacteria in the guts of people with higher wellbeing were, on average, better equipped to produce molecules with known neuroactive properties.

This study was significant for several reasons. It was large (over 2,000 participants across two cohorts). It was carefully controlled. It used both self-reported and clinician-diagnosed depression. It accounted for medication effects. And it went beyond simple taxonomic associations to examine functional capacity — what the bacteria were *capable* of doing, biochemically, in the brain-relevant domain. It remains the strongest population-level evidence linking specific gut bacteria to human mental health.

But — and this is a critical “but” — it was still an observational study. It could tell us that certain bacteria are depleted in depression. It could not tell us whether the depletion *causes* depression, *results from* depression, or reflects some third factor (diet, stress, medication history, genetics) that

causes both. The arrow of causation remains frustratingly ambiguous.

Clinical Trials: What the Evidence Actually Shows

The rubber meets the road in clinical trials — randomised, controlled experiments in which people with depression or anxiety are given either a probiotic or a placebo, and the outcomes are measured. There have now been dozens of such trials, and several meta-analyses have attempted to synthesise their results.

The overall picture is cautiously positive but far from conclusive. A 2025 meta-analysis pooling data from 23 randomised controlled trials found that probiotics produced a statistically significant reduction in depression scores, with a moderate-to-large effect size — and a smaller but still significant reduction in anxiety scores [REF:asad2025]. The strains most consistently associated with benefit were *Lactobacillus acidophilus*, *L. casei*, *L. plantarum*, *Bifidobacterium bifidum*, *B. lactis*, and *B. longum* — mostly the same genera that produce GABA and influence tryptophan metabolism in preclinical studies. *Bifidobacterium* strains appeared particularly effective for anxiety.

These numbers sound encouraging, and they are — with several important qualifications that any responsible discussion of this evidence must include.

The studies are small. Most trials enrolled fewer than 50 participants per arm. Small studies are prone to chance findings and to overestimating effect sizes — a well-known statistical phenomenon called the “winner’s curse,” in which the first studies to report a positive result tend to show larger effects than subsequent, larger replications.

The studies are heterogeneous. Different trials use different bacterial strains, different doses, different durations of treatment, different measures of depression and anxiety, and different patient populations (some recruit from the general population, others from clinical settings). This makes pooling results across studies statistically treacherous. The high heterogeneity in the meta-analyses (in some cases exceeding 90 per cent, meaning the studies are more different from each other than they are similar) suggests that the “average” effect may not represent any individual study particularly well.

Publication bias is likely. Trials that find a positive result are more likely to be published than trials that find no effect. This is a pervasive problem across all of clinical research, but it is particularly acute in emerging fields where the initial results are exciting and there is strong commercial interest. It is almost certain that negative probiotic trials exist in file drawers, unpublished.

The clinical significance is uncertain. A statistically significant reduction in a depression score is not the same as a clinically meaningful improvement in a person’s life. Many of the trials report improvements on standardised questionnaires (the Hamilton Depression Rating Scale, the Beck Depression Inventory) that are modest in absolute terms — a few points on a scale where a shift of several points is needed before a clinician would consider the patient improved. Whether participants *felt* meaningfully better is, in most trials, not reported.

None of this means psychobiotics do not work. It means we do not yet know, with the precision and confidence that evidence-based medicine demands, which strains work, for which patients, at what dose, and for how long. The field is where antidepressant pharmacology was in the 1960s — we can see that something is happening, but we have not yet figured out the prescribing. We are, in effect, trying to tune a radio when we have not yet mapped all the frequencies.

13.6 Autism Spectrum Disorder — Signal in the Noise?

The GI Connection

Of all the conditions linked to the gut-brain axis, none has attracted more public attention — or generated more controversy — than **autism spectrum disorder (ASD)**. The connection between ASD and the gut is not new to scientists, though it may surprise the general reader: gastrointestinal symptoms are strikingly common in autistic individuals. Depending on the study, between 30 and 70 per cent of children with ASD experience significant GI problems — constipation, diarrhoea, abdominal pain, bloating, and reflux — at rates roughly four times higher than in neurotypical children.

The GI symptoms are not incidental. Parents and clinicians have long observed that GI flare-ups in autistic children often coincide with worsening behavioural symptoms — increased irritability, repetitive behaviours, and social withdrawal. This temporal correlation, while not proof of a causal link, led researchers to ask an obvious question: is the gut microbiome different in autism, and if so, does it contribute to the condition?

The short answer is: the gut microbiome in ASD is different — but the differences are maddeningly inconsistent across studies, and the question of whether the microbiome contributes to ASD or is merely altered by it (through dietary preferences, restricted eating patterns, medication use, and stress) remains unresolved.

The Mouse Model: *Bacteroides fragilis* and Behaviour

The most elegant animal model study in the ASD-microbiome field came in 2013 from Elaine Hsiao (whom we met in section 13.2) and colleagues, working in Sarkis Mazmanian's laboratory at Caltech [REF:hsiao2013].

Hsiao used a mouse model of autism based on **maternal immune activation (MIA)** — a well-established paradigm in which pregnant mice are injected with a viral mimic (a molecule called polyinosinic:polycytidylic acid, or poly(I:C), which triggers an immune response without causing actual infection). The offspring of MIA mothers display a suite of behavioural abnormalities that parallel features of human autism: reduced social interaction, increased repetitive behaviours, altered communication (in mice, ultrasonic vocalisations), and increased anxiety. These MIA offspring also have a “leaky gut” — increased intestinal permeability — and an altered gut microbiome.

Hsiao's key experiment was to treat the MIA offspring with *Bacteroides fragilis* — the same immunomodulatory bacterium whose polysaccharide A drives regulatory T cell induction, as we discussed in Chapter 12. The results were striking. *B. fragilis* treatment corrected the intestinal permeability defect, restored a more normal microbial community, and — most remarkably — improved several of the autism-like behaviours. The treated mice communicated more, showed less anxiety, and engaged in fewer repetitive behaviours. Social behaviour, however, was not significantly rescued — a detail that is often omitted in popular accounts but is important for scientific accuracy.

The study also identified a candidate metabolite: **4-ethylphenylsulfate (4-EPS)**, a molecule produced by gut bacteria that was elevated in the blood of MIA offspring and restored to normal levels by *B. fragilis* treatment. When Hsiao injected 4-EPS alone into healthy mice, it induced anxiety-like behaviour — suggesting that the metabolite itself, rather than the bacterium, was at least partly responsible for the behavioural effects.

This was a beautifully designed study, and it remains one of the most cited papers in the gut-brain

field. But it is essential to recognise its limitations. The MIA mouse model captures some features of autism but not others. Autism in humans is genetically and phenotypically heterogeneous — a spectrum encompassing individuals with very different profiles of strengths and challenges. A mouse that is less social after its mother had an immune response during pregnancy is not an “autistic mouse” in any meaningful clinical sense. The study shows that a bacterium can influence behaviour through a gut-derived metabolite. It does not show that autism in humans is caused by, or treatable through, the gut microbiome.

Human Microbiome Studies: Inconsistency and Complexity

Dozens of studies have compared the gut microbiomes of autistic and neurotypical individuals. The results are, to put it charitably, inconsistent. Some studies find increased *Clostridium* species in ASD. Others do not. Some find decreased *Bifidobacterium*. Others find no difference. Some find altered *Bacteroides* levels. The direction of the alteration varies between studies. A 2019 study by Sharon and colleagues transplanted gut microbiota from human ASD donors into germ-free mice and found that the recipient mice developed autism-like behaviours — reduced social interaction, increased repetitive behaviour — with corresponding changes in brain gene expression and microbial metabolite profiles [REF:sharon2019]. This was provocative evidence, but the same caveats about germ-free mice and the limitations of behavioural assays apply.

The inconsistency across human studies likely reflects several factors. ASD is not one condition but many — a heterogeneous group of neurodevelopmental differences with varied genetic underpinnings, varied environmental contributions, and varied clinical presentations. Looking for “the autism microbiome” is a bit like looking for “the European climate” — the label covers such diverse territory that any single description is bound to miss most of the variation. There is no reason to expect a single autism microbiome any more than there is a single “autism gene.” Additionally, the diets of autistic children are often restricted (many children with ASD are highly selective eaters, with strong preferences for processed and low-fibre foods), and diet is the single largest determinant of gut microbiome composition. Separating the effects of the condition from the effects of the diet it produces is methodologically formidable.

Faecal Microbiota Transplantation: The Open-Label Promise and the RCT Reality

The most direct test of whether the microbiome can influence ASD symptoms in humans has come from **faecal microbiota transplantation (FMT)** trials — and the results illustrate, with painful clarity, the gap between preliminary excitement and rigorous evidence.

In 2017, Dae-Wook Kang, James Adams, and colleagues at Arizona State University published the results of an open-label trial — a study in which all participants receive the active treatment and both researchers and participants know what is being given — of a modified FMT protocol in 18 children with ASD and significant GI symptoms [REF:kang2017]. The protocol was intensive: two weeks of oral vancomycin (an antibiotic to suppress the existing microbiome), a bowel cleanse, and then seven to eight weeks of daily standardised human gut microbiota, delivered orally.

The results were encouraging. GI symptoms improved substantially — an approximately eighty per cent reduction in measures of constipation, diarrhoea, and abdominal pain. Autism-related behavioural symptoms also improved, by roughly twenty per cent on standardised assessment scales. The gut microbiome of the treated children became more diverse and showed increased abundance of *Bifidobacterium* and *Prevotella* — genera that had been depleted at baseline.

A two-year follow-up, published in 2019, reported that the improvements had not only persisted but continued to grow [REF:kang2019]. Professional evaluators, using standardised diagnostic instruments, found an average 45 per cent reduction in core ASD symptoms at the two-year mark compared with baseline. The microbiome changes — increased diversity and *Bifidobacterium/Prevotella* abundance — were still present.

These results generated enormous excitement, particularly among families of autistic children. But they need to be interpreted with caution, because the study had no control group. In an open-label trial, every participant knows they are receiving an active treatment. The placebo effect in behavioural interventions is large — particularly when the participants are children, the assessors are parents who desperately want improvement, and the treatment involves weeks of intensive medical attention. Without a placebo arm, it is impossible to know how much of the improvement was due to the microbiome transplant and how much was due to expectation, attention, and the natural variability of ASD symptoms over time.

The importance of this caveat was demonstrated in 2024, when Wan and colleagues published the results of a proper randomised, double-blind, placebo-controlled trial of oral FMT in 103 children with ASD [REF:wan2024]. The primary outcome — change in the Social Responsiveness Scale, one of the most widely used measures of autism-related social behaviour — showed no significant difference between the FMT group and the placebo group. There was a small improvement on a secondary measure, but it was modest and of uncertain clinical significance.

The contrast between the open-label results and the placebo-controlled results is sobering. It does not mean the microbiome has no role in ASD — the animal evidence is too strong to dismiss entirely, and the GI symptoms that are so common in autism clearly deserve treatment in their own right. But it does mean that the most direct test of the hypothesis — give a microbiome transplant and measure the effect on autism symptoms — has not yet produced the clear positive signal that many hoped for.

Navigating Hope and Evidence

The ASD-microbiome story carries a particular ethical weight. Autism is a condition around which enormous parental anxiety, commercial exploitation, and unproven therapies already circulate. Families of autistic children are, understandably, hungry for treatments that might help — and they are vulnerable to marketing that oversells preliminary findings. The history of autism is littered with “breakthrough” therapies — secretin injections, chelation therapy, hyperbaric oxygen — that generated excitement on the basis of small, uncontrolled studies and subsequently failed in rigorous trials.

The microbiome hypothesis is different from these in that it has genuine biological plausibility, supported by animal data and by the established connections between the gut, the immune system, and the brain. It is not pseudoscience. But it is also not yet a treatment — and presenting it as one, to families who are desperate for answers, risks causing harm: financial harm, if families pay for unproven FMT procedures; physical harm, if the procedures carry adverse effects (FMT is not risk-free — it involves transplanting a complex biological material that can contain pathogens); and psychological harm, if hopes are raised and then dashed.

The responsible position, as of 2026, is this: the gut-brain axis is real, the microbiome influences brain development and function through defined pathways, and it is plausible that microbiome-targeted interventions could, eventually, be part of a comprehensive approach to supporting aut-

istic individuals — particularly those with significant GI symptoms. But we are not there yet. The rigorous evidence is thin. The RCT results are disappointing. And the distance between a mouse that explores a maze differently and a child navigating the complexities of human social life is vast.

The Conversation Continues

The gut-brain axis is one of the most rapidly evolving areas of biomedical science, and it would be dishonest to pretend that the story is complete. It is not. We are, in many ways, at the beginning — we have identified the communication channels, mapped some of the molecules, and established that the microbiome can influence brain function in animal models with remarkable specificity. What we have not yet achieved is the translation of these findings into interventions that reliably improve human mental health.

The gap between the animal evidence and the clinical evidence is the central tension in this field. In mice, the effects of microbial manipulation on behaviour are large, consistent, and mechanistically tractable. In humans, the effects are small, inconsistent, and confounded by the complexity of human psychology, human genetics, and human lives. This is not unusual in biomedical research — the history of drug development is a long catalogue of therapies that worked spectacularly in mice and modestly (or not at all) in people. But it is a reminder that biological plausibility is not the same as clinical efficacy, and that the leap from mechanism to medicine is longer than it looks.

What is clear — and what the evidence does support, robustly — is that the gut and the brain are in continuous, bidirectional communication, and that the microbiome is a significant participant in that conversation. The vagus nerve carries microbial signals to the emotional brain. Bacterial metabolites influence the production of serotonin, GABA, and BDNF. The microbiome calibrates the stress response during a critical developmental window. Disruptions in the gut ecosystem produce changes in brain chemistry, immune function, and behaviour that are measurable, reproducible, and, in some cases, reversible.

In the next chapter, we follow this thread further — into the more speculative territory of whether microbes can influence not just how we feel but what we *want*. Chapter 14 examines the evidence that the microbiome may modulate food cravings, social behaviour, and even the decisions we think of as our own — and asks the uncomfortable question of who, exactly, is in charge.

Chapter 13 — Key References

- [REF:bravo2011] Bravo, J.A., Forsythe, P., Chew, M.V., Escaravage, E., Savignac, H.M., Dinan, T.G., Bienenstock, J. & Cryan, J.F. (2011). Ingestion of *Lactobacillus* strain regulates emotional behavior and central GABA receptor expression in a mouse via the vagus nerve. *Proceedings of the National Academy of Sciences*, 108(38), 16050–16055. DOI: 10.1073/pnas.1102999108
- [REF:yano2015] Yano, J.M., Yu, K., Donaldson, G.P., Shastri, G.G., Ann, P., Ma, L., Nagler, C.R., Ismagilov, R.F., Mazmanian, S.K. & Hsiao, E.Y. (2015). Indigenous bacteria from the gut microbiota regulate host serotonin biosynthesis. *Cell*, 161(2), 264–276. DOI: 10.1016/j.cell.2015.02.047

- [REF:strandwitz2019] Strandwitz, P., Kim, K.H., Terekhova, D., Liu, J.K., Sharma, A., Levering, J., McDonald, D., Dietrich, D., Ramadhar, T.R., Lekbua, A., Mroue, N., Liston, C., Stewart, E.J., Dubin, M.J., Zengler, K., Knight, R., Gilbert, J.A., Clardy, J. & Lewis, K. (2019). GABA-modulating bacteria of the human gut microbiota. *Nature Microbiology*, 4(3), 396–403. DOI: 10.1038/s41564-018-0307-3
- [REF:sudo2004] Sudo, N., Chida, Y., Aiba, Y., Sonoda, J., Oyama, N., Yu, X.N., Kubo, C. & Koga, Y. (2004). Postnatal microbial colonization programs the hypothalamic-pituitary-adrenal system for stress response in mice. *Journal of Physiology*, 558(Pt 1), 263–275. DOI: 10.1113/jphysiol.2004.063388
- [REF:bercik2011] Bercik, P., Denou, E., Collins, J., Jackson, W., Lu, J., Jury, J., Deng, Y., Blennerhassett, P., Macri, J., McCoy, K.D., Verdu, E.F. & Collins, S.M. (2011). The intestinal microbiota affect central levels of brain-derived neurotrophic factor and behavior in mice. *Gastroenterology*, 141(2), 599–609.e3. DOI: 10.1053/j.gastro.2011.04.052
- [REF:diazheijtz2011] Diaz Heijtz, R., Wang, S., Anuar, F., Qian, Y., Björkholm, B., Samuelsson, A., Hibberd, M.L., Forssberg, H. & Pettersson, S. (2011). Normal gut microbiota modulates brain development and behavior. *Proceedings of the National Academy of Sciences*, 108(7), 3047–3052. DOI: 10.1073/pnas.1010529108
- [REF:dinan2013] Dinan, T.G., Stanton, C. & Cryan, J.F. (2013). Psychobiotics: a novel class of psychotropic. *Biological Psychiatry*, 74(10), 720–726. DOI: 10.1016/j.biopsych.2013.05.001
- [REF:kelly2016] Kelly, J.R., Borre, Y., O'Brien, C., Patterson, E., El Aidy, S., Deane, J., Kennedy, P.J., Beers, S., Scott, K., Moloney, G., Hoban, A.E., Scott, L., Fitzgerald, P., Ross, P., Stanton, C., Clarke, G., Cryan, J.F. & Dinan, T.G. (2016). Transferring the blues: depression-associated gut microbiota induces neurobehavioural changes in the rat. *Journal of Psychiatric Research*, 82, 109–118. DOI: 10.1016/j.jpsychires.2016.07.019
- [REF:vallescolomer2019] Valles-Colomer, M., Falony, G., Darzi, Y., Tigchelaar, E.F., Wang, J., Tito, R.Y., Schiweck, C., Kurilshikov, A., Joossens, M., Wijmenga, C., Claes, S., Van Oudenhove, L., Zhernakova, A., Vieira-Silva, S. & Raes, J. (2019). The neuroactive potential of the human gut microbiota in quality of life and depression. *Nature Microbiology*, 4(4), 623–632. DOI: 10.1038/s41564-018-0337-x
- [REF:asad2025] Asad, A., Kirk, M., Zhu, S., et al. (2025). Effects of prebiotics and probiotics on symptoms of depression and anxiety in clinically diagnosed samples: systematic review and meta-analysis of randomized controlled trials. *Nutrition Reviews*, 83(7), e1504–e1520. DOI: 10.1093/nutrit/nuae177
- [REF:hsiao2013] Hsiao, E.Y., McBride, S.W., Hsien, S., Sharon, G., Hyde, E.R., McCue, T., Codelli, J.A., Chow, J., Reisman, S.E., Petrosino, J.F., Patterson, P.H. & Mazmanian, S.K. (2013). Microbiota modulate behavioral and physiological abnormalities associated with neurodevelopmental disorders. *Cell*, 155(7), 1451–1463. DOI: 10.1016/j.cell.2013.11.024
- [REF:sharon2019] Sharon, G., Cruz, N.J., Kang, D.W., Gandal, M.J., Wang, B., Kim, Y.M., Zink, E.M., Casey, C.P., Taylor, B.C., Lane, C.J., Bramer, L.M., Isern, N.G., Hoyt, D.W., Noecker, C., Sweredoski, M.J., Moradian, A., Borber, E., Janber, E.R., Knight, R., Geschwind, D.H., Mazmanian, S.K. et al. (2019). Human gut microbiota from autism spectrum disorder promote behavioral symptoms in mice. *Cell*, 177(6), 1600–1618.e17. DOI: 10.1016/j.cell.2019.05.004

- [REF:kang2017] Kang, D.W., Adams, J.B., Gregory, A.C., Borody, T., Chittick, L., Fasano, A., Khoruts, A., Geis, E., Maldonado, J., McDonough-Means, S., Pollard, E.L., Roux, S., Sadowsky, M.J., Lipson, K.S., Sullivan, M.B., Caporaso, J.G. & Krajmalnik-Brown, R. (2017). Microbiota Transfer Therapy alters gut ecosystem and improves gastrointestinal and autism symptoms: an open-label study. *Microbiome*, 5(1), 10. DOI: 10.1186/s40168-016-0225-7
- [REF:kang2019] Kang, D.W., Adams, J.B., Coleman, D.M., Pollard, E.L., Maldonado, J., McDonough-Means, S., Caporaso, J.G. & Krajmalnik-Brown, R. (2019). Long-term benefit of Microbiota Transfer Therapy on autism symptoms and gut microbiota. *Scientific Reports*, 9, 5821. DOI: 10.1038/s41598-019-42183-0
- [REF:wan2024] Wan, Y.P. et al. (2024). Effect of oral faecal microbiota transplantation intervention for children with autism spectrum disorder: a randomised, double-blind, placebo-controlled trial. *Clinical and Translational Medicine*, 14, e70006. DOI: 10.1002/ctm2.70006

Additional references to be verified and expanded during manuscript preparation.

Chapter 14: Behavioural Modulation — Are We the Puppet or the Puppeteer?

In the previous chapter, we traced the communication channels between the gut and the brain — the vagus nerve, the immune system, the endocrine system, and the metabolic pipeline of microbial molecules that, through various intermediaries, can influence mood, anxiety, and cognition. We established that the conversation is real, that the microbiome is a participant, and that disrupting the microbial community produces measurable changes in brain chemistry and behaviour.

Now we follow that thread into territory that is simultaneously more speculative and more unsettling. Because the question raised by the gut-brain axis is not just whether microbes can influence how we *feel*. It is whether they can influence what we *want* — what we eat, who we approach, how we behave. And if they can, it raises a question that most neuroscience textbooks do not attempt: who, exactly, is making the decisions?

This is also a chapter where we need to be particularly careful with the evidence. The popular press has seized on the idea that “your gut bacteria control your mind!” with an enthusiasm that far outstrips the science, producing headlines that flatten a nuanced and genuinely interesting story into a simple tale of microbial puppet masters. The reality, as we shall see, involves subtler mechanisms, more modest effects, and harder questions about where influence ends and control begins. Let us start with the most everyday example: the urge to eat.

14.1 Can Microbes Influence Food Cravings?

The Feedback Loop

You are standing in your kitchen at nine o'clock in the evening, and you want chocolate. Not a salad, not an apple, not the sensible leftover soup in the fridge — chocolate. The craving is specific, insistent, and hard to resist. You have always assumed this is your brain doing what brains do: seeking a reward, responding to stress, compensating for a long day. And you are not wrong. The brain's reward circuitry — the mesolimbic dopamine pathway, running from the ventral tegmental area to the nucleus accumbens — is heavily involved in food cravings, and it responds powerfully to sugar and fat.

But here is the question that the microbiome field has begun to ask: is the craving entirely yours? Or has it been, in some small way, suggested to you by the trillions of organisms living in your colon — organisms that have their own nutritional needs, their own metabolic preferences, and, as it turns out, their own chemical vocabulary for influencing the systems that determine what you eat?

The idea was laid out formally in 2014 by Joe Alcock, Carlo Maley, and Athena Aktipis in a paper in *BioEssays* that has become one of the most discussed theoretical frameworks in the microbiome field [REF:alcock2014]. Their argument was elegant and, at its core, evolutionary. Gut bacteria are not a homogeneous mass. They are a community of species with different metabolic specialisations, competing with each other for resources in the gut ecosystem. Some species thrive on dietary fibre

— the complex plant carbohydrates that human enzymes cannot digest. Others prefer simple sugars. Others metabolise amino acids, fats, or mucin (the glycoprotein that makes up the mucus layer of the gut). Each species does best when its preferred substrates are abundant in the gut lumen — which is to say, when the host eats the foods that feed that particular species.

Alcock and colleagues proposed that natural selection would favour microbes that could influence the host's dietary choices in their own favour. Not through conscious manipulation — bacteria do not have intentions — but through the chemical consequences of their metabolism. A bacterium that produces a metabolite which, as a side effect, makes the host feel slightly better after eating the food that feeds that bacterium would, over evolutionary time, outcompete a bacterium that did not. The mechanism need not be sophisticated. It need not involve a single “craving gene” in the bacterial genome. It only requires that microbial metabolites — produced for the bacterium's own metabolic reasons — happen to interact with the host's appetite-signalling pathways in a way that biases food choice.

This is not as far-fetched as it might sound. We already know, from the previous chapter, that gut bacteria produce neurotransmitters, short-chain fatty acids, and other neuroactive compounds that act on vagal afferents and enteroendocrine cells. We already know that these signals reach the brain's reward and appetite circuits. The question is whether these signals are strong enough, and specific enough, to influence what ends up on your plate.

The Machinery of Appetite

To understand how microbes might influence cravings, we need a brief tour of how appetite works — because appetite, it turns out, is not a single sensation. It is the integrated output of dozens of overlapping signals, and the gut microbiome has access to several of them.

When you eat, food enters the stomach and then the small intestine, where it is detected by specialised **enteroendocrine cells** — the same cells we met in Chapter 13 as producers of serotonin. These cells are scattered throughout the gut lining, and they function as chemical sensors, detecting nutrients in the gut lumen and responding by releasing hormones into the bloodstream and onto nearby nerve endings. The hormones they produce include **GLP-1** (glucagon-like peptide-1), which promotes satiety and slows gastric emptying; **PYY** (peptide YY), which reduces appetite and slows intestinal transit; and **CCK** (cholecystokinin), which signals fullness after a meal. Together, these hormones form the body's “I've had enough” system — a brake pedal that tells the brain to stop eating.

On the other side of the equation, when the stomach is empty and blood glucose is low, the stomach produces **ghrelin** — the “hunger hormone” — which travels to the hypothalamus and activates the **AgRP/NPY neurons** in the arcuate nucleus, a cluster of cells that drive the urge to seek food.

Now here is where the microbiome enters the picture. The short-chain fatty acids produced by bacterial fermentation of dietary fibre — acetate, propionate, and butyrate — are not just fuel for colonocytes and immune modulators. They also act on enteroendocrine cells. In 2021, Jiao and colleagues demonstrated in a mouse model that all three major SCFAs reduced body weight gain and appetite, but through different mechanisms: acetate increased GLP-1 and leptin secretion, propionate increased both PYY and GLP-1, and butyrate decreased overall feed intake [REF:jiao2021]. The same group had shown similar results in pigs the year before [REF:jiao2020]. The SCFAs were activating the “I've had enough” signalling cascade — not from the brain, but from the gut.

The mechanism involves a pair of receptors on enteroendocrine cells called **FFAR2** (free fatty acid

receptor 2) and **FFAR3** — G-protein coupled receptors that bind short-chain fatty acids and trigger hormone release. When SCFAs bind FFAR2 and FFAR3, the enteroendocrine cells release GLP-1 and PYY, which signal through the bloodstream and the vagus nerve to the hypothalamus, reducing appetite. A high-fibre diet, which produces abundant SCFAs through bacterial fermentation, thus generates a sustained satiety signal. A low-fibre diet — the processed, refined diet that dominates in industrialised countries — produces fewer SCFAs, a weaker satiety signal, and, potentially, greater hunger.

This is not “microbes controlling your appetite” in any dramatic sense. It is microbes producing metabolites, as a consequence of their own metabolism, that happen to feed into the host’s appetite-regulation system. But the net effect is the same: the composition of your gut microbiome, shaped by what you ate yesterday, influences how hungry you feel today and what your body signals you to eat tomorrow. It is a feedback loop. You feed the microbes. The microbes signal back.

A Bacterial Protein That Mimics a Satiety Hormone

The SCFA story is suggestive but indirect — the bacteria are not “trying” to modulate appetite; they are simply fermenting fibre, and the host’s signalling systems pick up the chemical traces. But one line of research has uncovered something more specific and more striking: a bacterial protein that is a structural mimic of a human satiety hormone.

The work comes from the laboratory of Sergey Fetissov, initially at Rouen University in France, and it centres on a protein called **ClpB** (caseinolytic protease B) — a protein produced by *Escherichia coli* and other members of the Enterobacteriaceae family [REF:fetissov2017].

ClpB is a chaperone protein — in bacteria, its job is to help refold proteins that have been damaged by heat or other stresses. It has nothing to do, from the bacterium’s perspective, with appetite. But Fetissov’s team discovered that a segment of ClpB shares structural similarity with **alpha-melanocyte-stimulating hormone (α-MSH)** — one of the key satiety signals in the human brain. α-MSH is produced by neurons in the hypothalamus, and it acts on the **melanocortin 4 receptor (MC4R)** to suppress appetite. Mutations in MC4R are the most common genetic cause of severe obesity in humans. α-MSH is, in the language of appetite neuroscience, one of the most powerful “stop eating” signals in the body.

ClpB, it turned out, could activate the same receptor. And because it is a protein — and therefore an antigen — the immune system mounts antibody responses against it. These anti-ClpB antibodies cross-react with α-MSH itself, because the two molecules look similar to the immune system. The result is a complex interplay between a bacterial protein, a human satiety hormone, and the immune system — all converging on appetite.

In 2020, Legrand and colleagues tested a ClpB-producing bacterium — *Hafnia alvei* strain HA4597 — in obese mice and found that it reduced body weight gain and food intake [REF:legrand2020]. A ClpB-deficient strain of *E. coli* had no such effect, confirming that the protein itself was responsible. The following year, Déchelotte and colleagues published the results of a randomised, double-blind, placebo-controlled trial in 236 overweight human subjects: those receiving *H. alvei* HA4597 were significantly more likely to achieve clinically meaningful weight loss, reported greater feelings of fullness, and showed greater reductions in hip circumference compared with placebo [REF:dechelotte2021].

This is remarkable. A bacterial protein, produced for the bacterium’s own housekeeping purposes, happens to resemble a human hormone closely enough to activate the same receptor and modulate

the same appetite circuit. Is this coincidence? Convergent evolution? Or has selection favoured bacteria that produce proteins which, by influencing host appetite in ways that increase the availability of the bacteria's preferred nutrients, give those bacteria a competitive advantage in the gut ecosystem?

Fetissov has argued for a version of the last hypothesis. In a 2016 review in *Nature Reviews Endocrinology*, he proposed that bacterial growth cycles — the lag phase, exponential growth phase, and stationary phase that all bacterial populations go through after a meal — are synchronised with host appetite cycles, and that the metabolites and proteins produced at each phase feed back to the host's hunger and satiety systems [REF:fetissov2016]. In this model, bacteria do not “decide” to make the host hungry. But the chemical consequences of their growth dynamics — including the production of ClpB during the stationary phase, when the bacteria have finished dividing and nutrients are running low — happen to coincide with the host's transition from hunger to satiety. The bacterium's life cycle and the host's meal cycle are, in effect, entrained — locked together by chemical signals that neither party consciously controls.

You Are What Your Microbiome Eats

There is one more piece of this puzzle, and it is perhaps the most thought-provoking: the relationship between dietary patterns and microbial community structure is not one-directional. Your diet shapes your microbiome — this has been demonstrated repeatedly and convincingly. David and colleagues showed in 2014 that switching between a plant-based and an animal-based diet altered the gut microbiome within a single day [REF:david2014]. People who eat high-fibre, plant-rich diets tend to harbour **Prevotella**-dominated communities, which are specialised for fermenting complex carbohydrates. People who eat high-fat, high-protein Western diets tend to harbour **Bacteroides**-dominated communities, which are better equipped to metabolise animal fats and amino acids.

But the microbiome also shapes your dietary preferences — or at least creates the conditions under which certain preferences are reinforced. A *Prevotella*-dominated gut, thriving on fibre, produces abundant SCFAs, which signal satiety and reward through the pathways described above. A person with that microbiome may find that a plant-rich meal leaves them feeling satisfied and well. If they switch to a low-fibre diet, the *Prevotella* population crashes, SCFA production drops, the satiety signals weaken, and they may feel less satisfied after meals — hungrier, more likely to snack, more drawn to energy-dense processed foods that provide quick reward without requiring microbial fermentation. The *Bacteroides* that expand to fill the niche have their own metabolic preferences and produce their own signalling molecules, which may reinforce the new dietary pattern.

This is what ecologists call a **positive feedback loop** — or, less charitably, a trap. Your diet shapes your microbiome, and your microbiome shapes the chemical environment that influences your next dietary choice. Eat fibre, grow fibre-fermenters, feel satisfied on fibre. Eat sugar, grow sugar-lovers, feel unsatisfied without sugar. The loop is not deterministic — you can override it with conscious effort, and dietary changes can shift the microbial community within days. But it means that the first few days of a new diet are the hardest, not just because of willpower, but because the microbial community has not yet caught up with the new plan. The old residents are still sending the old signals.

In 2023, Kim and colleagues at an animal model study demonstrated a striking version of this feedback loop [REF:kim2023]. They colonised germ-free rats with microbiota from donors who had been eating either a high-fat or a standard diet. The rats receiving high-fat-diet microbiota showed

reduced dopamine levels and dopamine receptor (DRD2) expression in the nucleus accumbens — the brain’s reward centre — and reduced motivation for palatable food. When the researchers severed the vagal afferents, the dopamine deficits were restored. The microbiome was altering the reward value of food, through the vagus nerve, in a way that was consistent with the dietary history of the donor — not the recipient. The craving, in a sense, had been transplanted along with the bacteria.

This is not mind control. But it is a form of influence that operates below the threshold of conscious awareness, mediated by chemical signals that the host cannot perceive as “coming from the microbiome” and has no reason to question. When you crave chocolate at nine in the evening, the craving feels like *yours*. And in every meaningful sense, it is — you are the one who experiences it, you are the one who acts on it, and you are the one who can choose to eat the soup instead. But the intensity of that craving, and the specific character of the reward you experience when you satisfy it, may be shaped in part by the metabolic preferences of organisms you have never seen and will never meet.

14.2 Social Behaviour and Microbial Signalling in Animal Models

The Bacterium That Made Mice Social

In Chapter 13, we saw that germ-free mice display profound social deficits — reduced interaction with strangers, diminished curiosity about unfamiliar animals, and abnormal social investigation. We noted that these deficits could be partially corrected by microbial colonisation, and that specific metabolites like hippuric acid appeared to act through the oxytocin system. Now we look more closely at the social effects of the microbiome — because the story is both more specific and more surprising than a general “microbes are needed for normal behaviour” narrative would suggest.

The most striking experiment in this domain came in 2016 from Shelly Buffington, Mauro Costa-Mattioli, and colleagues at Baylor College of Medicine, published in *Cell* [REF:buffington2016]. The experiment began with a dietary intervention, not a microbial one. Buffington’s team fed pregnant mice a high-fat diet — the kind of diet that, in humans, is associated with obesity and, increasingly, with neurodevelopmental differences in offspring. The pups born to high-fat-diet mothers were, by standard laboratory measures, socially impaired. They spent less time interacting with unfamiliar mice, showed reduced preference for social novelty, and displayed a suite of behaviours that, in mouse terms, look like the social deficits seen in autism spectrum models.

When the researchers examined the gut microbiome of these socially impaired pups, they found it was different from that of pups born to mothers on a normal diet — less diverse, with particular species depleted. One species stood out by its absence: ***Lactobacillus reuteri***, a bacterium we have encountered before (in Chapter 3, as an early coloniser of the infant gut, and in Chapter 13, in the context of GABA production and vagal signalling).

Buffington then asked a simple question: would restoring the missing bacterium restore the missing behaviour? She added *L. reuteri* to the drinking water of the socially impaired pups. The answer was yes — and the specificity of the rescue was remarkable. The *L. reuteri*-treated mice showed normalised social behaviour: they spent more time with unfamiliar mice, showed greater social preference, and behaved indistinguishably from pups born to normal-diet mothers on social interaction tests. A single bacterial species had corrected a social deficit caused by a maternal dietary perturbation.

The mechanism, as Buffington's team and subsequent studies by Sgritta and colleagues in 2018 elucidated, ran through the vagus nerve to the brain's oxytocin system [REF:sgritta2018]. *L. reuteri* treatment increased the number of **oxytocin-producing neurons** in the **paraventricular nucleus (PVN)** of the hypothalamus — the brain region that serves as the body's master regulator of oxytocin release. Oxytocin, as we noted briefly in the previous chapter, is the neuropeptide most closely associated with social bonding, trust, and social recognition. It is released during positive social interactions, during breastfeeding, during physical contact, and — in the experimental system that Buffington and Costa-Mattioli used — during social exploration of unfamiliar mice.

The *L. reuteri* effect on social behaviour required two things: an intact vagus nerve and functional oxytocin receptors. When the researchers severed the vagus nerve, the social rescue disappeared — just as Bravo's *L. rhamnosus* effect on anxiety (Chapter 13) had disappeared after vagotomy. When they blocked oxytocin receptors pharmacologically, the social rescue also disappeared. The pathway was clear: *L. reuteri* in the gut □ vagal signalling □ oxytocin neurons in the hypothalamus □ social behaviour. A single species, acting through a single nerve, modulating a single neuropeptide, rescuing a complex behaviour.

Sgritta's 2018 follow-up went further. It showed that *L. reuteri* could rescue social deficits in multiple mouse models — not just the high-fat-diet model, but also genetic models of autism (the Shank3 knockout) and the valproic acid model (in which prenatal exposure to the anti-epileptic drug produces autism-like features in offspring). Across all three models, the bacterium restored social behaviour through the same vagus-oxytocin pathway. And critically, the rescue was *specific to social behaviour*. *L. reuteri* did not correct anxiety, repetitive behaviours, or other non-social features of the models. It did one thing — restored sociality — through one pathway, with one molecular target.

This specificity is both scientifically important and narratively irresistible. It is hard not to marvel at a world in which a bacterium in the gut, producing metabolites that travel up a nerve to a specific cluster of neurons in the hypothalamus, can make a mouse more interested in other mice. But we must, as always, temper the marvel with caution. These are mouse experiments. The social “deficit” is defined by how long a mouse spends sniffing a stranger in a cage, not by the nuanced social difficulties that characterise human autism. And *L. reuteri* is not a treatment for autism — a point we will return to.

Chemical Calling Cards: Microbes and Social Communication

The idea that microbes influence social behaviour is not limited to the gut-brain axis. In some species, microbes are the social behaviour — or at least, they produce the chemical signals that make social communication possible.

In 2013, Kevin Theis and colleagues at Michigan State University published a study on the scent-marking behaviour of hyenas — the chemical signals that these animals deposit from specialised glands near the anus to communicate identity, sex, reproductive status, and social rank to other members of the group [REF:theis2013]. Hyena scent marks are complex chemical cocktails, and the signals they carry vary between individuals and between social groups. Theis's team showed that the volatile compounds in hyena scent marks — the molecules that give each mark its distinctive smell — are not produced by the hyena itself. They are produced by **bacteria living in the scent glands**. The bacterial communities in the scent glands varied between species (spotted hyenas versus striped hyenas), between sexes, and between reproductive states — and the volatile fatty acid profiles they produced varied in parallel. The hyena's social identity card was, in effect, written by

its bacteria.

This is an example of what biologists call the **fermentation hypothesis** of mammalian chemical communication: the idea that many of the scent signals used by mammals for social recognition are not synthesised by the animal's own cells but are produced by symbiotic bacteria metabolising the secretions of specialised glands. The animal provides the substrate; the bacteria provide the signal. If this hypothesis is correct — and the evidence from hyenas, badgers, and several other species supports it — then microbial communities are not merely passengers in social behaviour. They are the chemical infrastructure on which social communication is built.

An even more striking example comes from a different kingdom entirely. In 2010, Gil Sharon and colleagues at Tel Aviv University demonstrated that the gut bacteria of *Drosophila melanogaster* — the common fruit fly — could influence mating preferences [REF:sharon2010]. Flies raised on different diets developed different gut bacterial communities, and these flies preferentially mated with partners that had been raised on the same diet. The mating preference was not genetic. It was microbial. When the flies were treated with antibiotics — wiping out their gut bacteria — the mating preference disappeared. When the bacteria were restored, it returned. The flies were choosing mates based, at least in part, on microbial cues — probably volatile compounds produced by the bacteria that altered the flies' scent profile and, consequently, their attractiveness to potential partners.

This experiment made headlines, and justifiably so — it suggested that the microbiome could influence one of the most fundamental behaviours in biology: mate choice. But it also attracted criticism. Subsequent replication attempts by other laboratories produced inconsistent results, and the precise mechanism — which bacteria, which volatile compounds, which sensory receptors in the fly — has not been fully resolved. The finding remains provocative rather than established, a signpost pointing toward a possibility rather than a proven mechanism.

What These Studies Tell Us

The animal studies on microbes and social behaviour converge on a consistent theme: the gut microbiome, through the vagus nerve and the oxytocin system, can influence the neural circuits that regulate social motivation, social recognition, and social bonding. The effects are real, reproducible in at least some model systems, and mechanistically tractable — we can trace the pathway from bacterium to metabolite to nerve to brain region to behaviour with reasonable precision.

But the studies also reveal the limits of what we know. All of the strongest evidence comes from rodent models. The gap between a mouse that sniffs a stranger more enthusiastically after receiving *L. reuteri* and a human who navigates the complexities of friendship, empathy, romantic attachment, and social hierarchy is not a gap that a single bacterial species can bridge. Human social behaviour is shaped by culture, language, experience, personality, and a cerebral cortex that is vastly more complex than anything a mouse possesses. The microbiome may contribute a note to the chord, but the chord has many notes, and the human brain is a far more complex instrument than the mouse brain.

14.3 The Evolutionary Logic: Microbial Manipulation for Transmission Advantage

The Puppet Masters of the Animal Kingdom

Before we assess the strength of the evidence that our *own* microbiome manipulates our behaviour, it is worth pausing to acknowledge that in the wider biological world, microbial manipulation of host behaviour is not a hypothesis. It is an established fact. Some of the most spectacular examples in all of biology involve parasites and pathogens that hijack their hosts' nervous systems with terrifying precision.

The most famous example — the one that has launched a thousand popular science articles — is **Toxoplasma gondii**, a single-celled parasite that can infect virtually any warm-blooded animal but can only sexually reproduce in the gut of a cat. *Toxoplasma* faces an evolutionary problem: it needs to get from its intermediate host (typically a rodent) into a cat's stomach. The most efficient way to do that is to make the rodent *more likely to be eaten by a cat*.

In 2000, Manuel Berdoy, Joanne Webster, and David Macdonald at the University of Oxford published the experiment that demonstrated this with devastating clarity [REF:berdoy2000]. They infected rats with *Toxoplasma* and tested their response to cat urine — a stimulus that uninfected rats avoid with an innate, deep-seated aversion that does not require learning. Rats do not need to encounter a cat to know that cat smell means danger; the fear is hardwired.

In the infected rats, this aversion was not merely reduced. It was *reversed*. The *Toxoplasma*-infected rats were attracted to cat urine. They approached it, lingered near it, and showed no signs of the fear response that would normally keep them alive. Other aspects of their behaviour were preserved — they were still cautious about other predator odours, still normally reactive to non-predator stimuli. The change was specific: the parasite had selectively disabled the fear of the one predator that it needed the rat to encounter.

The mechanism, as subsequent research has revealed, involves *Toxoplasma*'s remarkable trick of producing its own **dopamine**. The parasite genome encodes two copies of **tyrosine hydroxylase** — the rate-limiting enzyme in dopamine synthesis — and the tissue cysts that *Toxoplasma* forms in the brain are preferentially located in brain regions associated with fear and reward, including the amygdala [REF:mconkey2013]. The parasite is, in effect, rewiring the host's emotional circuitry from the inside, using the host's own neurotransmitter language.

Joanne Webster, who has spent much of her career studying *Toxoplasma* manipulation, refers to this as “fatal attraction” — a term she has explored in a comprehensive 2025 review [REF:webster2025]. The attraction is fatal for the rat (which gets eaten) and beneficial for the parasite (which completes its life cycle). It is manipulation in the truest sense: a change in host behaviour that benefits the parasite at the host's expense, achieved through a specific molecular mechanism that natural selection has refined over millions of years of coevolution.

Toxoplasma is not alone. The **rabies virus** modifies the behaviour of infected animals in a way that maximises its transmission through saliva — infected animals become aggressive, tend to bite, and produce copious amounts of virus-laden saliva. The virus has essentially weaponised the host's behaviour to serve its own dispersal. The zombie-ant fungus **Ophiocordyceps** takes behavioural manipulation to perhaps its most baroque extreme. When an ant is infected, the fungus takes control of its motor behaviour, compelling it to climb to a specific height on a plant, bite down on the underside of a leaf with a “death grip,” and die there — in the precise location that is optimal for

the fungus to sporulate and rain spores down on the colony below [REF:debekker2014]. Different *Ophiocordyceps* species manipulate different ant species, and de Bekker and colleagues at Penn State showed in 2014 that the fungus produces different arrays of metabolites in response to different host brains — it is, in a sense, custom-tuning its chemical manipulation to the specific neural architecture of each host species.

These are extraordinary examples. They demonstrate beyond any doubt that microorganisms can evolve sophisticated mechanisms for manipulating host behaviour in service of their own transmission. But they are also, critically, examples of **parasites with clear transmission strategies** — organisms that need to get from one host to another and can increase their chances by modifying what the current host does. The question for the human microbiome is different, and more subtle.

The Commensal's Dilemma

The gut microbiome is not *Toxoplasma*. It is not a single-species parasite with a defined life cycle that requires the host to be eaten. It is a community of hundreds of species, most of which are commensals or mutualists — organisms that benefit from the relationship without harming the host, or that benefit both parties. Their “transmission strategy,” such as it is, involves being passed from mother to child during birth and breastfeeding, spread through household contact, and potentially transmitted through food and environmental exposure. They do not need the host to be eaten by a predator. They need the host to *stay alive*, to *eat well*, and to *interact with other humans* — because social contact is one of the primary routes through which gut bacteria spread between individuals.

This distinction matters enormously for the evolutionary logic of behavioural manipulation. In 2018, Kevin Johnson and Kevin Foster published a rigorous analysis in *Nature Reviews Microbiology* that addressed the question directly: “Why does the microbiome affect behaviour?” [REF:johnson2018]. They distinguished three possible explanations for any observed microbial effect on host behaviour:

Manipulation: the microbe benefits from the behavioural change at the host's expense. This is what *Toxoplasma* does. It requires that the behavioural change increase the microbe's fitness (reproduction or transmission) and that natural selection can act on the responsible genes.

Byproduct: the behavioural change is an incidental side effect of microbial metabolism that is not under selection. The bacterium produces a neurotransmitter for its own metabolic reasons, the host's nervous system happens to detect it, and behaviour changes — but the change is not adaptive for the bacterium. It is noise, not signal.

Mutualism: the behavioural change benefits both the microbe and the host. A bacterium that makes the host more social, for example, might benefit both parties — the host gains the advantages of social living, and the bacterium gains more opportunities for horizontal transmission to new hosts.

Johnson and Foster argued — and this is the critical point — that true manipulation by the gut microbiome faces a severe evolutionary obstacle. For manipulation to evolve, the manipulating genes must be under positive selection — they must increase in frequency because the organisms carrying them have higher fitness. In a single-species parasite like *Toxoplasma*, this is straightforward: the parasite is a genetically coherent unit, and any gene that increases transmission success will be favoured. But the gut microbiome is a *community* of hundreds of genetically unrelated species. A manipulation gene in one species would benefit that species — but the fitness cost to the host would

be shared across the entire community. This is the biological equivalent of the “tragedy of the commons”: the benefit of manipulation is private (captured by one species), but the cost is public (borne by all species, including the manipulator, if the host’s health declines).

Moreover, the gut microbiome is genetically diverse and rapidly evolving. Any manipulation strategy that required coordination among multiple species — “let’s all produce the same craving signal to make the host eat more sugar” — would be vulnerable to cheaters: species that benefit from the host’s behaviour change without contributing to the costly signalling. Such cooperative manipulation is, on evolutionary grounds, implausible in a diverse microbial community.

This does not mean the gut microbiome has *no* effect on behaviour. It clearly does — the evidence from the previous chapter and this one is too strong to dismiss. But Johnson and Foster’s analysis suggests that most of these effects are best understood as **byproducts** of microbial metabolism rather than **evolved manipulation strategies**. The bacteria are not trying to make you eat chocolate. They are fermenting whatever you give them, producing metabolites as a consequence of their own biochemistry, and some of those metabolites happen to interact with the host’s nervous system. The effect on behaviour is real, but it is not designed. It is an emergent property of a complex system, not a purpose-built tool.

Jaroslav Flegr and the Human Implications of *Toxoplasma*

There is, however, one provocative exception to the general rule that behavioural manipulation by the human microbiome remains unproven — and it involves *Toxoplasma gondii* itself.

Toxoplasma is not, technically, part of the microbiome. It is a parasite, and infection (toxoplasmosis) is a disease state. But roughly a third of the world’s human population carries chronic *Toxoplasma* infection — acquired from undercooked meat, contaminated water, or contact with cat faeces — and in most of these individuals, the infection is asymptomatic. The parasite forms cysts in the brain that persist for life, quietly producing dopamine and other neuroactive compounds.

Jaroslav Flegr, an evolutionary biologist at Charles University in Prague, has spent more than thirty years studying the behavioural effects of chronic *Toxoplasma* infection in humans — and his findings, while controversial, are consistent and thought-provoking. In large population studies, Flegr has found that *Toxoplasma*-positive individuals differ from uninfected controls on several personality measures: they score differently on tests of conscientiousness, impulsivity, and risk-taking, and the effects are sex-specific — infected men tend to become less conscientious and more risk-prone, while infected women show a different pattern [REF:flegr2023]. In a 2024 study of 557 young adults, Flegr and colleagues found that *Toxoplasma* infection was associated with lower IQ scores in men, with the effects modulated by sex and Rh blood factor [REF:flegr2024].

These findings are correlational, not causal — it is possible that the personality differences precede infection (perhaps more impulsive people are more likely to be exposed to *Toxoplasma* through risky dietary habits) rather than result from it. And the effect sizes are modest. No one who meets a *Toxoplasma*-positive person at a dinner party would notice anything unusual. But the consistency of the findings across multiple studies, combined with the known mechanism (dopamine production in brain tissue), makes this the strongest candidate for genuine parasitic behavioural manipulation in humans.

It is worth noting what Flegr himself emphasises: the behavioural changes in humans are far subtler than the dramatic “fatal attraction” seen in rats. This makes evolutionary sense. The cat-predation pathway that *Toxoplasma* exploits in rodents does not apply to humans — we are not regularly eaten

by cats. Any behavioural effects in humans are likely spillover from mechanisms that evolved for the rodent host, not adaptations for human transmission. We are, in the language of parasitology, an **accidental host** — infected, but not the intended target of the manipulation.

14.4 Free Will, Determinism, and the Uncomfortable Philosophical Implications

The Question You've Been Waiting For

Let us be honest: the reason the microbiome-and-behaviour story captures the public imagination is not the short-chain fatty acids, the vagal signalling, or the GABA receptors. It is the philosophical implication. If microbes can influence what we eat, how we feel, who we approach, and how we respond to stress — if our behaviour is shaped, even in part, by organisms that are not us — then what does that mean for the concept of free will? Are we, as the title of this chapter asks, the puppet or the puppeteer?

The question is seductive, but it is also, in an important sense, not new. Neuroscience has been eroding the folk concept of free will for decades. We know that brain tumours can change personality. We know that a mutation in the MAO-A gene alters impulse control. We know that the microenvironment of the prefrontal cortex during adolescent development shapes adult decision-making. We know that unconscious neural processes precede conscious decisions by measurable intervals — the famous Libet experiments of the 1980s showed that the brain begins preparing a movement several hundred milliseconds before the person becomes aware of “deciding” to move. If all of this is true — if behaviour is shaped by genetics, brain chemistry, developmental history, and neural processing that precedes conscious awareness — then the microbiome is just one more influence in a long list of influences that the conscious self does not control and, for the most part, does not even perceive.

The philosopher Daniel Dennett would recognise this situation. Dennett argued that free will is not about escaping causation — about somehow stepping outside the chain of physical events that produces behaviour. It is about having the kind of causal architecture that allows flexible, considered, context-appropriate responses. A thermostat responds to temperature, but it has no free will. A chess computer responds to the board position with enormous sophistication, but most people would not call its moves “free.” A human being, integrating information from sensory systems, memory, social context, emotional state, and — yes — gut-derived chemical signals, produces behaviour that is responsive to reasons, amenable to deliberation, and capable of being modified by new information. That is what free will means in any useful sense.

The microbiome does not change this picture. It adds one more input to an already complex system. Your gut bacteria may nudge you toward the chocolate rather than the salad. But you can still, if you choose, eat the salad. The nudge is real — it is mediated by specific chemical signals acting on specific neural circuits — but it is a *nudge*, not a command. It is more like the background music in a restaurant (which studies show influences how much people eat and what they order) than like a puppeteer pulling strings. It shapes the probability landscape of your choices without determining any individual choice.

The Real Philosophical Challenge

The interesting philosophical question raised by the microbiome is not about free will. It is about **identity**.

When we say “I want chocolate,” who is the “I”? In the traditional model, the “I” is the conscious self — the person, the mind, the unified agent that makes choices. But the human body is, as we have been arguing throughout this book, not a unified agent. It is a community — a collection of human cells, bacterial cells, fungal cells, and viral genomes that have been coevolving for hundreds of millions of years. The human genome contains about 20,000 protein-coding genes. The gut microbiome contributes at least 2 million more. The metabolic capabilities of the microbiome dwarf those of the human genome. The immune system, as we saw in Chapter 12, is shaped and calibrated by microbial signals. The nervous system, as we have seen in this chapter and the last, receives microbial inputs that influence mood, motivation, and social behaviour.

If the microbiome contributes to the physiological processes that produce behaviour — and it demonstrably does — then the “I” that experiences a craving or an emotion is not a purely human “I.” It is a holobiont “I” — the integrated output of a multi-species organism that includes bacteria, fungi, and viruses as functional components. The chocolate craving is not coming from “outside” you, from alien organisms imposing their will on a helpless host. It is coming from *within* you, from a system that includes both human and microbial components, because the system that is “you” has always included both.

This may sound like a semantic trick, but it has real implications for how we think about health, responsibility, and intervention. If your microbiome contributes to your food preferences, then a diet change is not just an exercise of willpower against desire. It is an ecological intervention — a restructuring of the microbial community that generates the desires. If your microbiome contributes to your stress response, then managing stress is not just a psychological problem. It is a microbiological one. The boundaries between “me” and “my microbiome” are not the hard borders we might imagine. They are negotiated, permeable, and — like everything else in biology — the product of a long evolutionary conversation.

14.5 How Seriously Should We Take the “Microbes Control Behaviour” Hypothesis?

A Sober Assessment

We have covered a great deal of territory in this chapter — from chocolate cravings to zombie ants, from oxytocin-producing bacteria to parasites that rewrite the fear response. It is time to step back and ask, honestly, how much of this applies to everyday human life.

The answer requires distinguishing several claims that are often conflated:

Claim 1: The gut microbiome produces molecules that interact with host neural signalling. This is established beyond reasonable doubt. The evidence from Chapters 12, 13, and this chapter is overwhelming. Gut bacteria produce neurotransmitters, SCFAs, tryptophan metabolites, and proteins like ClpB that interact with appetite, mood, and stress-response pathways. The communication channels — vagal, immune, endocrine, metabolic — are well characterised. This is not controversial.

Claim 2: Changes in the gut microbiome are associated with changes in behaviour in animal models. This is also well established. Germ-free mice have abnormal stress responses, social deficits, and altered anxiety. Specific bacterial species can rescue specific behavioural deficits through defined molecular pathways. Faecal transplants can transfer behavioural traits between mouse strains. The animal evidence is strong.

Claim 3: The gut microbiome meaningfully influences human behaviour in everyday life. This is plausible but largely unproven. The human evidence consists of correlational studies (people with certain microbial profiles have different mental health outcomes), small and heterogeneous clinical trials (probiotics show modest effects on depression and anxiety scores), and mechanistic reasoning extrapolated from animal models. The gap between “a single bacterium changes GABA receptor expression in a mouse brain via the vagus nerve” and “my gut bacteria make me crave pizza” has not been bridged with direct experimental evidence in humans.

Claim 4: Gut microbes have evolved to manipulate human behaviour for their own benefit. This is the weakest claim, and the one that generates the most exciting headlines. As Johnson and Foster’s analysis makes clear, the evolutionary conditions for true manipulation by commensal gut bacteria are stringent and largely unmet. Most microbial effects on behaviour are better explained as byproducts of microbial metabolism or as mutualistic effects that benefit both host and microbe. The genuine cases of behavioural manipulation in nature — *Toxoplasma*, *Ophiocordyceps*, rabies — involve single-species parasites with clear transmission incentives, not diverse commensal communities.

Where This Leaves Us

The microbiome does not control your behaviour. But it contributes to the physiological substrate from which behaviour emerges. It is one voice in a very large choir — a choir that includes your genes, your developmental history, your hormones, your immune system, your sleep patterns, your social environment, and the breakfast you did or did not eat this morning. The microbiome’s contribution is real, measurable, and — in principle — modifiable. But it is modest, context-dependent, and mediated through pathways that the conscious mind can override.

The popular narrative of “gut bacteria control your mind” is misleading not because it is entirely wrong, but because it misrepresents the nature of biological causation. In a complex system, no single factor “controls” anything. Everything influences everything else, through feedback loops and regulatory networks that make linear cause-and-effect thinking inadequate. The microbiome is part of the system. It is not the master of it.

If there is a practical takeaway from this chapter, it is this: the things that are good for your microbiome — a fibre-rich diet, regular physical activity, adequate sleep, reduced chronic stress, and limited unnecessary antibiotic use — are the same things that are good for your brain. This is not a coincidence. It is a consequence of the fact that these systems are connected, that they evolved together, and that what disrupts one tends to disrupt the other. You cannot optimise your brain while ignoring your gut, any more than you can tune a piano by adjusting only the keys and ignoring the strings.

In the next chapter, we leave the brain and turn to a different conversation between the microbiome and the body — one conducted not through nerves but through hormones. Chapter 15 explores the microbial influence on metabolism, the endocrine system, and the body’s internal clocks.

Chapter 14 — Key References

- [REF:alcock2014] Alcock, J., Maley, C.C. & Aktipis, C.A. (2014). Is eating behavior manipulated by the gastrointestinal microbiota? Evolutionary pressures and potential mechanisms. *BioEssays*, 36(10), 940–949. DOI: 10.1002/bies.201400071
- [REF:fetissov2016] Fetissov, S.O. (2017). Role of the gut microbiota in host appetite control: bacterial growth to animal feeding behaviour. *Nature Reviews Endocrinology*, 13(1), 11–25. DOI: 10.1038/nrendo.2016.150
- [REF:fetissov2017] Fetissov, S.O. (2017). Involvement of gut bacteria in appetite control. *Biologie Aujourd'hui*, 211(1), 29–37. DOI: 10.1051/jbio/2017007
- [REF:legrand2020] Legrand, R., Lucas, N., Dominique, M., Azhar, S., Deroissart, C., Le Solliec, M.-A., Rondeaux, J., Nobis, S., Guérin, C., Léon, F., do Rego, J.-C., Pons, N., Le Chatelier, E., Ehrlich, S.D., Lambert, G. & Déchelotte, P. (2020). Commensal *Hafnia alvei* strain reduces food intake and fat mass in obese mice — a new potential probiotic for appetite and body weight management. *International Journal of Obesity*, 44(5), 1041–1051. DOI: 10.1038/s41366-019-0515-9
- [REF:dechelotte2021] Déchelotte, P., Breton, J., Trotin-Piccolo, C., Grube, B., Couturier, C., Legrand, R., Bothe, G., Fetissov, S.O. & Lambert, G. (2021). The probiotic strain *Hafnia alvei* HA4597 improves weight loss in overweight subjects under moderate hypocaloric diet: a proof-of-concept, multicenter randomized, double-blind placebo-controlled study. *Nutrients*, 13(6), 1902. DOI: 10.3390/nu13061902
- [REF:jiao2021] Jiao, A., Yu, B., He, J., Yu, J., Zheng, P., Luo, Y., Luo, J., Yan, H., Wang, Q., Wang, H., Mao, X. & Chen, D. (2021). Sodium acetate, propionate, and butyrate reduce fat accumulation in mice via modulating appetite and relevant genes. *Nutrition*, 87–88, 111198. DOI: 10.1016/j.nut.2021.111198
- [REF:jiao2020] Jiao, A., Yu, B., He, J., Yu, J., Zheng, P., Luo, Y., Luo, J., Mao, X. & Chen, D. (2020). Short chain fatty acids could prevent fat deposition in pigs via regulating related hormones and genes. *Food & Function*, 11(2), 1845–1855. DOI: 10.1039/c9fo02585e
- [REF:kim2023] Kim, J.S., Kirkland, R.A., Lee, S.H., Cawthon, C.R., Rzepka, K.W., Minaya, D.M., de Lartigue, G., Bhatt, A.P., Carroll, I.M. & de La Serre, C.B. (2023). The gut-brain axis mediates bacterial driven modulation of reward signaling. *Molecular Metabolism*, 75, 101764. DOI: 10.1016/j.molmet.2023.101764
- [REF:david2014] David, L.A., Maurice, C.F., Carmody, R.N., Gootenberg, D.B., Button, J.E., Wolfe, B.E., Ling, A.V., Devlin, A.S., Varma, Y., Fischbach, M.A., Biddinger, S.B., Dutton, R.J. & Turnbaugh, P.J. (2014). Diet rapidly and reproducibly alters the human gut microbiome. *Nature*, 505(7484), 559–563. DOI: 10.1038/nature12820
- [REF:buffington2016] Buffington, S.A., Di Prisco, G.V., Auchtung, T.A., Ajami, N.J., Petrosino, J.F. & Costa-Mattioli, M. (2016). Microbial reconstitution reverses maternal diet-induced social and synaptic deficits in offspring. *Cell*, 165(7), 1762–1775. DOI: 10.1016/j.cell.2016.06.001
- [REF:sgritta2018] Sgritta, M., Dooling, S.W., Buffington, S.A., Momin, E.N., Francis, M.B., Britton, R.A. & Costa-Mattioli, M. (2019). Mechanisms underlying microbial-mediated

changes in social behavior in mouse models of autism spectrum disorder. *Neuron*, 101(2), 246–259.e6. DOI: 10.1016/j.neuron.2018.11.018

- [REF:theis2013] Theis, K.R., Venkataraman, A., Dycus, J.A., Koonter, K.D., Schmitt-Matzen, E.N., Wagner, A.P., Holekamp, K.E. & Schmidt, T.M. (2013). Symbiotic bacteria appear to mediate hyena social odors. *Proceedings of the National Academy of Sciences*, 110(49), 19832–19837. DOI: 10.1073/pnas.1306477110
- [REF:sharon2010] Sharon, G., Segal, D., Ringo, J.M., Hefetz, A., Zilber-Rosenberg, I. & Rosenberg, E. (2010). Commensal bacteria play a role in mating preference of *Drosophila melanogaster*. *Proceedings of the National Academy of Sciences*, 107(46), 20051–20056. DOI: 10.1073/pnas.1009906107
- [REF:berdoy2000] Berdoy, M., Webster, J.P. & Macdonald, D.W. (2000). Fatal attraction in rats infected with *Toxoplasma gondii*. *Proceedings of the Royal Society B*, 267(1452), 1591–1594. DOI: 10.1098/rspb.2000.1182
- [REF:mcconkey2013] McConkey, G.A., Martin, H.L., Bristow, G.C. & Webster, J.P. (2013). *Toxoplasma gondii* infection and behaviour — location, location, location? *Journal of Experimental Biology*, 216(1), 113–119. DOI: 10.1242/jeb.074153
- [REF:webster2025] Webster, J.P. (2025). Who makes the decisions? Uncovering the evolutionary implications and clinical applications of *Toxoplasma gondii*'s Fatal Feline Attraction. *Folia Parasitologica*, 72. DOI: 10.14411/fp.2025.016
- [REF:flegr2023] Flegr, J., Ullmann, J. & Toman, J. (2023). Parasitic manipulation or side effects? The effects of past *Toxoplasma gondii* and *Borrelia* spp. infections on human personality and cognitive performance are not mediated by impaired health. *Folia Parasitologica*, 70. DOI: 10.14411/fp.2023.020
- [REF:flegr2024] Flegr, J., Chvátalová, V., Příplatová, L. & Toman, J. (2024). Cognitive effects of *Toxoplasma* and CMV infections: a cross-sectional study of 557 young adults considering modulation by sex and Rh factor. *Pathogens*, 13(5), 363. DOI: 10.3390/pathogens13050363
- [REF:johnson2018] Johnson, K.V.-A. & Foster, K.R. (2018). Why does the microbiome affect behaviour? *Nature Reviews Microbiology*, 16(10), 647–655. DOI: 10.1038/s41579-018-0014-3
- [REF:debekker2014] de Bekker, C., Quevillon, L.E., Smith, P.B., Fleming, K.R., Ghosh, D., Patterson, A.D. & Hughes, D.P. (2014). Species-specific ant brain manipulation by a specialized fungal parasite. *BMC Evolutionary Biology*, 14, 166. DOI: 10.1186/s12862-014-0166-3

Additional references to be verified and expanded during manuscript preparation.

Chapter 15: Metabolism, Hormones, and the Endocrine Microbiome

In the previous two chapters, we followed the conversation between the gut microbiome and the brain — a conversation conducted through nerves, neurotransmitters, and the subtle chemical nudges that may shape what we eat, how we feel, and who we seek out. We established that the microbiome is not merely a digestive workforce but a participant in the body's internal communications, producing molecules that interact with neural circuits in ways that are real, measurable, and — depending on your philosophical temperament — either fascinating or mildly unsettling.

Now we turn to a different communication network — one that is, in many ways, even more pervasive. The endocrine system.

If the nervous system is the body's telephone network — fast, targeted, point-to-point — then the endocrine system is its postal service. Hormones are chemical messages released into the bloodstream by specialised glands, carried throughout the body, and read by any cell that has the right receptor. The process is slower than neural signalling (minutes to hours rather than milliseconds), but it is broad, persistent, and powerful. Hormones regulate metabolism, growth, reproduction, stress, sleep, appetite, bone density, and mood. They are the body's long-range regulators — the thermostat settings that determine how fast you burn energy, how you store fat, when you feel hungry, how your body handles sugar, and whether your reproductive system is switched on or off.

And the microbiome, it turns out, has its fingers in all of it.

This should not surprise us, given what we have learned so far. We have already seen that the gut microbiome produces neurotransmitters, calibrates the immune system, and modulates the stress hormone axis. But the endocrine story extends well beyond the HPA axis we explored in Chapter 13. The microbiome participates in bile acid signalling — a hormonal conversation between the gut and the liver that regulates metabolism, fat storage, and blood sugar. It modulates the body's pool of circulating oestrogens — with implications for breast cancer, endometriosis, and bone health. It influences thyroid function — the metabolic thermostat that sets the pace of virtually every cell in the body. And it appears to interact with the body's circadian clock — the internal timekeeping system that tells your cells what time of day it is and adjusts their behaviour accordingly.

This chapter is about those conversations. It is also, necessarily, about the limits of what we know — because the endocrine microbiome is a younger field than the gut-brain axis, and many of its claims rest on animal models and correlational studies that have not yet been confirmed in rigorous human trials. We will flag the evidence quality as we go, as we have throughout this book. But the picture that is emerging — of a microbiome that participates in hormonal regulation at a systemic level — is, even in outline, remarkable.

15.1 Bile Acid Metabolism and the Gut-Liver Axis

A Molecule With Two Jobs

To understand how the microbiome talks to the liver — and, through it, to the entire metabolism — we need to start with bile acids. You have probably encountered bile, at least in passing, as a yellowish-green fluid associated with digestion. If you have ever vomited on an empty stomach, the bitter, acrid liquid that came up was bile. It is not a pleasant introduction to one of the most important signalling molecules in the body.

Bile acids are synthesised in the liver from cholesterol — yes, the same cholesterol that has spent decades as a dietary villain, although its reputation is more nuanced than the headlines suggest. The liver produces about 500 milligrams of bile acids per day, packages them with other components (phospholipids, bilirubin, water) into bile, and stores the mixture in the gallbladder. After a meal, the gallbladder contracts and squirts bile into the upper small intestine — the duodenum — where it performs its first and most obvious job: emulsifying dietary fats.

Think of bile acids as biological detergent. Fats do not dissolve in water, and the contents of the small intestine are essentially aqueous. Without bile acids, dietary fats would clump together in large, indigestible globules, inaccessible to the lipase enzymes that need to break them down. Bile acids surround these fat globules and break them into tiny droplets — a process called emulsification — increasing the surface area available for enzymatic digestion by several hundred-fold. This is the textbook job of bile acids, the one taught in every introductory physiology course, and it is genuinely important. Without it, you cannot absorb fat-soluble vitamins (A, D, E, K) or digest most dietary lipids.

But bile acids have a second job — one that was discovered much more recently and is, for our purposes, far more interesting. Bile acids are hormones. Or, more precisely, they are signalling molecules that activate specific receptors in cells throughout the body, triggering cascading changes in gene expression that regulate metabolism, inflammation, and energy balance. This second job was essentially unknown before the late 1990s, and its full implications are still being worked out. But it has transformed our understanding of bile acids from inert digestive detergents into potent metabolic regulators — and it is the doorway through which the gut microbiome enters the story.

The Recycling Loop

Before we get to the signalling, we need to understand the geography. Bile acids do not make a one-way trip from the liver to the toilet. They are far too valuable for that. After being secreted into the small intestine, approximately 95 per cent of bile acids are reabsorbed in the final segment of the small intestine — the **ileum** — through a dedicated transporter protein. From there, they travel via the portal vein back to the liver, which recycles them into new bile. This circuit — liver to gut to liver — is called the **enterohepatic circulation**, and a typical bile acid molecule completes it six to eight times per day. The total bile acid pool in the body is only about 3 to 5 grams, but because it recycles so efficiently, it processes roughly 20 to 30 grams of fat per day. It is a remarkably economical system.

The 5 per cent that escapes reabsorption — the fraction that makes it past the ileum and into the colon — is where the microbiome takes over.

The Microbial Transformation

The bile acids that the liver produces are called **primary bile acids**. In humans, the two main primary bile acids are **cholic acid** and **chenodeoxycholic acid**. Before being secreted into bile, the liver conjugates them — attaches either glycine or taurine to them, producing molecules like glycocholic acid and taurocholic acid. This conjugation makes them more water-soluble and more effective as detergents.

When these conjugated primary bile acids reach the colon — the 5 per cent that was not reabsorbed — they encounter the densest microbial community in the body. And the bacteria get to work.

The first thing the gut bacteria do is **deconjugate** the bile acids — strip off the glycine or taurine tag that the liver attached. This is performed by an enzyme called **bile salt hydrolase (BSH)**, which is produced by a remarkably wide range of gut bacteria, including species of *Lactobacillus*, *Bifidobacterium*, *Clostridium*, *Bacteroides*, and *Enterococcus*. BSH activity is so common in the gut microbiome that it has been proposed as a core function — something that virtually every gut microbial community can do, regardless of its exact species composition [REF:wahlstrom2016].

The second thing the bacteria do is more transformative. Through a series of enzymatic reactions — principally **7 α -dehydroxylation**, performed mainly by *Clostridium* species in cluster XIVa — they convert the primary bile acids into **secondary bile acids**. Cholic acid becomes **deoxycholic acid**. Chenodeoxycholic acid becomes **lithocholic acid**. These secondary bile acids have different chemical properties, different signalling activities, and different effects on host metabolism than their primary precursors.

The transformation matters because primary and secondary bile acids activate different receptors with different affinities — and therefore trigger different downstream effects. The microbiome, by controlling the ratio of primary to secondary bile acids, is effectively tuning the hormonal signal.

The Receptors: FXR and TGR5

The two most important bile acid receptors are the **farnesoid X receptor (FXR)** and the **Takeda G-protein-coupled receptor 5 (TGR5)**, sometimes called GPBAR1.

FXR is a nuclear receptor — meaning it sits inside the cell, in the nucleus, and directly regulates gene expression when activated. It is expressed in the liver, the intestine, the kidneys, and the adrenal glands. When bile acids bind to FXR, they trigger a cascade of effects: reduced bile acid synthesis in the liver (a negative feedback loop), improved insulin sensitivity, reduced fat synthesis, reduced inflammation, and altered glucose metabolism. FXR activation is, broadly speaking, metabolically beneficial — it promotes the kind of metabolic housekeeping that keeps blood sugar stable, lipid levels in check, and the liver healthy.

In 2013, Sama Sayin and colleagues at the University of Gothenburg published a landmark study that revealed how the microbiome modulates FXR signalling [REF:sayin2013]. Working with germ-free mice, they profiled bile acid composition throughout the gut and liver and discovered something unexpected. Germ-free mice had elevated levels of a bile acid called **tauro- β -muricholic acid (T β MCA)** — a conjugated primary bile acid that the liver produces but that gut bacteria normally break down. T β MCA, it turned out, is a natural **antagonist** of FXR — it blocks the receptor, preventing it from being activated by other bile acids. In conventionally raised mice, gut bacteria eliminated T β MCA through deconjugation and transformation, allowing FXR to function normally. In germ-free mice, T β MCA accumulated, FXR was suppressed, and the downstream metabolic con-

sequences were substantial: increased bile acid synthesis, altered cholesterol metabolism, and disrupted signalling of fibroblast growth factor 15 (FGF15) — a hormone produced in the ileum under FXR control that regulates bile acid production in the liver.

The study demonstrated, with molecular precision, that the gut microbiome does not merely process bile acids passively. It determines which bile acids are present, and by doing so, it controls whether a critical metabolic receptor is activated or suppressed. The microbiome was not just digesting leftovers. It was editing the hormonal mail.

TGR5 is a membrane receptor — it sits on the cell surface rather than in the nucleus, and it responds particularly strongly to secondary bile acids, especially lithocholic acid and deoxycholic acid. TGR5 activation in muscle and brown adipose tissue increases energy expenditure by promoting the conversion of inactive thyroid hormone (T4) into active thyroid hormone (T3) — a connection we will return to in section 15.3. TGR5 activation in the intestine stimulates the release of **GLP-1** — the satiety hormone we encountered in Chapter 14, which is also the target of the blockbuster diabetes and obesity drugs semaglutide and tirzepatide. TGR5 activation in immune cells — particularly macrophages — suppresses the production of pro-inflammatory cytokines, reducing inflammation.

The secondary bile acids that activate TGR5 most strongly are the ones produced by gut bacteria. Without microbial transformation, primary bile acids dominate, TGR5 activation is weaker, and the downstream effects — energy expenditure, GLP-1 release, anti-inflammatory signalling — are diminished. The microbiome, through its bile acid chemistry, is tuning both receptors simultaneously: removing an FXR antagonist (by breaking down T β MCA) and producing TGR5 agonists (by generating secondary bile acids). Both adjustments push host metabolism in the same direction — toward better glucose regulation, more efficient energy use, and reduced inflammation.

The Feedback Loop: Bile Acids Shape the Microbiome Too

The conversation between bile acids and the microbiome is not one-directional. Bile acids are, among other things, antimicrobial. They disrupt bacterial cell membranes, and different bile acids are toxic to different bacterial species. The composition of the bile acid pool therefore shapes which bacteria can survive in the gut — and the bacteria that survive, in turn, determine which bile acids are produced. It is a feedback loop, similar in principle to the dietary feedback loops we discussed in Chapter 14.

Wahlström and colleagues, in an influential 2016 review in *Cell Metabolism*, described this as a “crosstalk” — a bidirectional conversation in which bile acids and gut bacteria continuously adjust each other’s composition [REF:wahlstrom2016]. FXR activation in the small intestine, for example, induces the production of antimicrobial peptides that selectively suppress certain bacterial species while sparing others. The surviving bacteria produce different secondary bile acids, which alter FXR and TGR5 signalling, which changes antimicrobial peptide production, which reshapes the bacterial community again. The system is dynamic, self-regulating, and exquisitely sensitive to perturbation.

This matters clinically because disruptions to the bile acid–microbiome conversation are increasingly implicated in metabolic disease. In obesity, the gut microbiome is altered in ways that shift bile acid profiles toward compositions that favour fat storage and insulin resistance. In non-alcoholic fatty liver disease (NAFLD) — now the most common liver disease in the Western world, affecting roughly a quarter of the global population — bile acid signalling is disrupted, FXR activation is reduced, and the liver accumulates fat that it cannot efficiently export or burn. In type 2 diabetes,

altered bile acid profiles contribute to the impaired glucose regulation that defines the disease.

The therapeutic implications are being actively explored. FXR agonists — drugs that activate the receptor pharmacologically — have shown promise in clinical trials for NAFLD and liver fibrosis. **Obeticholic acid**, a synthetic bile acid and potent FXR agonist, was the first such drug to reach advanced clinical trials, though its path to approval has been complicated by side effects, including increased LDL cholesterol. More recently, researchers have begun asking whether modifying the gut microbiome — through diet, probiotics, or targeted interventions — might be a less blunt instrument for tuning bile acid signalling. The logic is straightforward: if the microbiome controls which bile acids are present, then reshaping the microbiome could reshape the hormonal signal without the off-target effects of a pharmaceutical hammer.

Whether this will work in practice remains to be seen. But the principle is clear: the gut-liver axis, mediated by bile acids, is a hormonal conversation in which the microbiome is not an eavesdropper but a co-author.

15.2 The Estrobolome — Microbial Regulation of Oestrogen

The Hormone That Needed a Recycling Plant

If bile acids are the metabolic postal service, oestrogen is the body's long-range contractor for a bewildering range of projects. Oestrogen — or, more precisely, the family of oestrogens, of which **oestradiol (E2)** is the most potent — is best known for its role in female reproductive biology: regulating the menstrual cycle, supporting pregnancy, and maintaining the health of the uterus, breasts, and vaginal tissue. But oestrogen's job description extends far beyond reproduction. It maintains bone density (which is why osteoporosis rates soar after menopause, when oestrogen levels drop). It protects cardiovascular health (which is why premenopausal women have significantly lower rates of heart disease than age-matched men). It modulates the immune system, influences brain function, and regulates aspects of lipid and glucose metabolism. Oestrogen receptors are found in virtually every tissue in the body — including the gut, the liver, the brain, bone, fat, and muscle.

Given the breadth of its influence, anything that alters circulating oestrogen levels has the potential to affect the entire body. And the gut microbiome, it turns out, is one of the most important regulators of how much oestrogen circulates in the blood.

Enterohepatic Circulation — Again

The mechanism involves a familiar concept: enterohepatic circulation. We just met it in the context of bile acids, and the oestrogen story runs through the same plumbing.

After oestrogens are produced by the ovaries (and, to a lesser extent, by fat tissue, the adrenal glands, and — in men — the testes), they circulate through the bloodstream and exert their effects on target tissues. Eventually, they arrive at the liver, which inactivates them through a process called **conjugation** — attaching a chemical group (usually glucuronic acid, in a reaction called glucuronidation) that renders the oestrogen water-soluble and biologically inactive. The conjugated oestrogens are then excreted in bile, travel to the intestine, and — if nothing intervenes — continue through the colon and are eliminated in the faeces.

But something does intervene. The gut microbiome.

Certain gut bacteria produce an enzyme called **β -glucuronidase**, which does exactly what you might guess from the name: it removes the glucuronic acid tag that the liver attached. This deconjugation reactivates the oestrogen, converting it from an inert, water-soluble waste product back into a biologically active hormone. The reactivated oestrogen is then reabsorbed through the intestinal wall, re-enters the bloodstream, and resumes its hormonal activity. The liver tried to excrete it. The gut bacteria sent it back.

In 2011, Claudia Plottel and Martin Blaser — the same Martin Blaser whose “disappearing microbiota hypothesis” we encountered in Chapter 3 — gave this process a name. They called the collection of gut microbial genes capable of metabolising oestrogens the **estrobolome** [REF:plottel2011]. The term was modelled on “microbiome” and “metabolome,” and it captured a simple but powerful idea: the gut microbiome, through its β -glucuronidase activity, functions as a regulator of circulating oestrogen levels. A microbiome with high β -glucuronidase activity sends more oestrogen back into circulation. A microbiome with low activity allows more oestrogen to be excreted. The net effect is that the composition of the gut microbiome influences the total body burden of circulating oestrogen — and, through it, every oestrogen-dependent process in the body.

The Evidence

The first direct evidence linking gut microbial diversity to circulating oestrogen levels in humans came from a 2012 cross-sectional study by Roberto Flores, James Goedert, and colleagues at the National Cancer Institute [REF:flores2012]. The study measured faecal microbiome composition, faecal β -glucuronidase and β -glucosidase activity, and urinary oestrogen levels in 51 adult volunteers (25 men, 7 postmenopausal women, and 19 premenopausal women).

The results were striking — and sex-specific. In men and postmenopausal women, total urinary oestrogens were strongly and directly associated with gut microbiome richness and diversity. Participants with more diverse gut microbiomes had higher circulating oestrogen levels. The correlation coefficients were remarkably strong for this type of study — $R \geq 0.50$, with p -values ≤ 0.003 . In premenopausal women, no such association was found — which makes biological sense, because in premenopausal women, the dominant source of oestrogen is the ovaries, which produce such large quantities that the microbiome’s contribution to the circulating pool is proportionally small. After menopause, when ovarian production ceases, the estrobolome becomes a relatively more important determinant of whatever oestrogen remains in circulation.

The finding was correlational, not causal — the study could not prove that gut diversity *caused* the higher oestrogen levels, as opposed to some third factor causing both. But it was consistent with the estrobolome hypothesis, and it has since been supported by additional studies linking gut microbial β -glucuronidase activity to oestrogen metabolism in both animal models and human cohorts.

Breast Cancer: The Uncomfortable Implication

The estrobolome matters clinically because oestrogen is a double-edged hormone. The same molecule that protects bones and hearts also fuels certain cancers. Approximately 70 to 80 per cent of breast cancers are **oestrogen-receptor-positive (ER+)** — meaning the cancer cells carry oestrogen receptors and their growth is stimulated by oestrogen. This is why anti-oestrogen therapies — tamoxifen, aromatase inhibitors — are the backbone of treatment for ER+ breast cancer. Anything that increases circulating oestrogen levels, or increases the duration of oestrogen exposure over a lifetime, is a risk factor for ER+ breast cancer. This is why early menarche, late menopause,

nulliparity (never having been pregnant), and postmenopausal hormone replacement therapy all increase breast cancer risk — they all increase the total lifetime exposure to oestrogen.

If the estrobolome modulates circulating oestrogen levels — and the evidence from Flores and others suggests it does — then the composition of the gut microbiome is, at least in principle, a modifiable risk factor for ER+ breast cancer. A microbiome with high β -glucuronidase activity, which recycles more oestrogen back into circulation, could increase cancer risk. A microbiome with lower activity, which allows more oestrogen to be excreted, could reduce it.

A 2025 review by Larnder and colleagues in the *International Journal of Cancer* examined the evidence for this connection and found it suggestive but still incomplete [REF:larnder2025]. Several case-control studies have reported altered gut microbiome composition in women with breast cancer compared with healthy controls, with reduced diversity and shifts in specific bacterial taxa. Some studies have found elevated β -glucuronidase activity in breast cancer patients. But the review also highlighted significant limitations: most studies are small, cross-sectional (making it impossible to determine whether the microbial changes preceded or followed the cancer), and narrowly focused on β -glucuronidase — ignoring other microbial pathways that may also metabolise oestrogens.

The honest assessment is that the estrobolome–breast cancer hypothesis is biologically plausible, consistent with the available evidence, and not yet proven. It is a hypothesis that deserves larger, prospective studies — studies that follow women over time, measure their gut microbiome composition before cancer develops, and determine whether microbial profiles predict subsequent cancer risk. Such studies are underway, but results are years away.

Beyond Breast Cancer: Oestrogen, Bones, and the Rest

The estrobolome's influence is not limited to cancer risk. If gut bacteria modulate circulating oestrogen, they potentially influence every oestrogen-dependent process — and the list is long.

Osteoporosis. The dramatic bone loss that follows menopause is driven primarily by oestrogen withdrawal. Oestrogen inhibits osteoclasts (the cells that break down bone) and supports osteoblasts (the cells that build it). When oestrogen drops, bone resorption outpaces formation, and bones become progressively thinner and more fragile. If the estrobolome contributes to postmenopausal oestrogen levels — even modestly — then gut microbiome composition could influence the rate of bone loss. Animal studies support this: germ-free mice have altered bone metabolism, and certain probiotic interventions have been shown to reduce bone loss in ovariectomised mice (a standard model for postmenopausal osteoporosis). Human evidence is early-stage but consistent with the hypothesis.

Endometriosis. This painful condition, in which tissue resembling the uterine lining grows outside the uterus, is oestrogen-dependent — the ectopic tissue is stimulated by oestrogen, which is why hormonal suppression is a mainstay of treatment. Several studies have reported altered gut microbiome composition in women with endometriosis, including elevated β -glucuronidase-producing bacteria. Whether this is a cause, a consequence, or a coincidence is not yet clear.

Cardiovascular disease. The cardiovascular protection that premenopausal women enjoy relative to men is widely attributed to oestrogen's beneficial effects on blood vessel function, lipid profiles, and inflammation. The sharp increase in cardiovascular risk after menopause tracks the decline in oestrogen. If the estrobolome contributes to residual postmenopausal oestrogen levels, it could influence cardiovascular risk — though this remains speculative.

Male health. The estrobolome is not exclusively a female story. Men produce oestrogen too — in smaller quantities, mainly through the conversion of testosterone to oestradiol by the enzyme aromatase in fat tissue. Oestrogen plays important roles in male bone health, brain function, and cardiovascular health. The gut microbiome's β -glucuronidase activity operates in men just as it does in women, recycling conjugated oestrogens back into the circulation. The Flores study found the same association between gut diversity and urinary oestrogens in men as in postmenopausal women. The implications are less well studied, but the biology is there.

The Testobolome: A Newer Story

The estrobolome concept has recently been extended to other sex hormones. In 2025, researchers proposed the term **testobolome** to describe the gut microbial pathways that metabolise testosterone and other androgens [REF:li2025]. Several gut bacterial species, including members of *Clostridium* and *Ruminococcus*, have been shown to possess enzymes capable of modifying testosterone — hydroxylating it, reducing it, or converting it to other androgenic compounds. The clinical implications are still emerging, but the principle is the same: the gut microbiome can modify sex hormones, and by doing so, it can influence the hormonal environment of the entire body.

The estrobolome — and its newer sibling, the testobolome — represent a conceptual shift. They reframe the gut microbiome not just as a digestive partner or an immune educator, but as an endocrine organ in its own right — one that participates in the regulation of sex hormones with consequences that ripple through bone, breast, brain, heart, and reproductive tissue. The shift is not yet fully proven, and the clinical applications are not yet clear. But the direction is unmistakable: the microbiome is part of the hormonal system, and we are only beginning to understand what that means.

15.3 Thyroid Function and the Microbiome

The Master Thermostat

If oestrogen is the body's long-range contractor, thyroid hormones are its thermostat. The thyroid gland — a small, butterfly-shaped organ at the front of the neck — produces hormones that regulate the metabolic rate of virtually every cell in the body. When thyroid hormone levels are adequate, metabolism hums along normally: you maintain a stable body temperature, your heart beats at an appropriate rate, your gut moves food through at a reasonable speed, your brain functions clearly, and your energy levels are steady. When thyroid hormone is insufficient — a condition called **hypothyroidism** — everything slows down. You feel cold, tired, constipated, mentally foggy, and your weight creeps up despite no change in diet. When there is too much — **hyperthyroidism** — everything speeds up: rapid heartbeat, weight loss, anxiety, tremor, diarrhoea, and intolerance of heat.

Thyroid disorders are strikingly common. Hypothyroidism affects roughly 5 per cent of the population in developed countries, with women affected five to eight times more often than men. The most common cause, in regions with adequate iodine intake, is **Hashimoto's thyroiditis** — an autoimmune condition in which the immune system attacks the thyroid gland, gradually destroying its ability to produce hormones. Given what we learned in Chapter 12 about the microbiome's role in immune calibration — and particularly its influence on the balance between inflammatory

and regulatory immune responses — it should not be surprising that researchers have begun asking whether the gut microbiome plays a role in thyroid disease.

The Gut-Thyroid Axis: Mechanisms

The relationship between the gut microbiome and the thyroid operates through several interconnected mechanisms — some well established, others still under investigation.

Nutrient absorption. The thyroid gland requires specific micronutrients to function: **iodine** (the essential building block of thyroid hormones — each molecule of the main thyroid hormone, thyroxine or T₄, contains four iodine atoms), **selenium** (required by the deiodinase enzymes that convert the inactive T₄ into the active form, triiodothyronine or T₃), **iron** (required by thyroid peroxidase, the enzyme that catalyses the incorporation of iodine into thyroid hormone), and **zinc** (involved in thyroid hormone synthesis and receptor binding). The gut microbiome influences the absorption of all of these micronutrients. Dysbiosis can impair absorption — through damage to the intestinal lining, competition for nutrients, or alteration of the gut environment — potentially starving the thyroid of the raw materials it needs [REF:knezevic2020].

This might sound like a minor contribution, but consider selenium. The deiodinase enzymes are **selenoproteins** — they contain the amino acid selenocysteine at their active site and cannot function without it. There are three deiodinases in the human body: type 1 (DIO1), which operates mainly in the liver and kidney; type 2 (DIO2), which operates in the brain, pituitary, and brown adipose tissue; and type 3 (DIO3), which inactivates thyroid hormones. Together, they control the local concentration of active thyroid hormone in every tissue. If selenium absorption is impaired by an altered gut microbiome, deiodinase activity drops, T₄-to-T₃ conversion slows, and the net effect is a functional hypothyroidism — even if the thyroid gland itself is producing adequate amounts of T₄. The thermostat is set correctly, but the heating system cannot convert the fuel into warmth.

Immune modulation. Hashimoto's thyroiditis is an autoimmune disease, and autoimmune diseases are, as we discussed in Chapter 12, intimately connected to the gut microbiome. The mechanisms we described — molecular mimicry, epithelial barrier dysfunction, and the balance between pro-inflammatory Th17 cells and regulatory T cells — all apply here. Multiple studies have reported altered gut microbiome composition in patients with Hashimoto's thyroiditis: reduced diversity, decreased abundance of *Bifidobacterium* and butyrate-producing *Clostridium* clusters, and increased abundance of potentially inflammatory genera including *Bacteroides*, *Klebsiella*, and *Proteobacteria* [REF:shi2024]. A 2024 review in *Frontiers in Cellular and Infection Microbiology* summarised the evidence across autoimmune, neoplastic, and nodular thyroid conditions and found consistent — if imperfectly replicated — associations between gut dysbiosis and thyroid autoimmunity [REF:sun2024].

The immune connection also runs through the phenomenon of **molecular mimicry** — the possibility that bacterial proteins in the gut resemble thyroid antigens closely enough to trigger a cross-reactive immune response. This has been proposed for Hashimoto's but, as with many molecular mimicry hypotheses, the direct evidence in humans remains circumstantial rather than conclusive.

The intestinal barrier. The “leaky gut” phenomenon — increased intestinal permeability allowing bacterial products to enter the bloodstream — is relevant here for the same reasons it is relevant to other autoimmune conditions. Increased permeability exposes the immune system to microbial antigens that it would not normally encounter, potentially triggering or exacerbating autoimmune responses. Thyroid hormones themselves influence gut motility and intestinal barrier function, cre-

ating another bidirectional feedback loop: hypothyroidism slows gut motility, which can promote bacterial overgrowth and dysbiosis, which further impairs barrier function, which worsens immune activation, which damages the thyroid further. The cycle is familiar from the stress-microbiome vicious cycle we described in Chapter 13. The endocrine system and the microbiome can trap each other in mutually reinforcing dysfunction.

Direct microbial effects on thyroid hormones. Perhaps the most intriguing — and least well understood — mechanism is the possibility that gut bacteria can directly metabolise thyroid hormones. Some bacterial species possess deiodinase-like enzymes and can convert T₄ to T₃, or can bind and sequester iodothyronines in the gut lumen. The gut has been estimated to serve as a reservoir for roughly 20 per cent of the body's total T₃, and microbial activity in the colon may contribute to the regulation of this pool. A 2025 review in *Frontiers in Microbiology* described this as one of the “pathogenic mechanisms” of gut dysbiosis in hypothyroidism — the idea that an altered microbiome could directly deplete or redistribute thyroid hormones in the gut, contributing to the systemic hormone deficiency [REF:chen2025].

Thyroid Medication and the Microbiome

There is a practical clinical dimension to the gut-thyroid connection that directly affects millions of patients. **Levothyroxine** — synthetic T₄ — is one of the most prescribed medications in the world (consistently among the top three most prescribed drugs in the United States, with roughly 80 to 100 million prescriptions per year). It is the standard treatment for hypothyroidism, and getting the dose right is notoriously tricky. Patients often require dose adjustments, and their hormone levels can fluctuate despite taking the same dose — a frustration familiar to anyone who has managed hypothyroidism, either as a patient or as a clinician.

One reason for this variability may be the gut microbiome. Levothyroxine is absorbed in the small intestine, and absorption efficiency depends on the integrity of the intestinal lining, the pH of the gut environment, and the presence of competing substances. Gut dysbiosis — particularly conditions that alter intestinal permeability or promote small intestinal bacterial overgrowth (SIBO) — can impair levothyroxine absorption, requiring higher doses to achieve the same blood levels. Conversely, successful treatment of SIBO in hypothyroid patients has been reported to improve levothyroxine absorption, sometimes necessitating a dose reduction.

This has led some researchers to propose that thyroid patients should be evaluated for gut health as part of their clinical management — not as an alternative to standard thyroid treatment, but as a way of optimising it. If a patient's levothyroxine dose requirement is unexpectedly high, or if their thyroid levels are erratic despite good compliance, considering the microbiome as a confounding factor is not unreasonable.

The IMITHOT Trial: Testing the Hypothesis

The most ambitious direct test of the gut-thyroid hypothesis is the **IMITHOT trial** — a double-blinded, randomised controlled trial being conducted in the Netherlands to evaluate whether faecal microbiota transplantation (FMT) can improve thyroid function in patients with subclinical autoimmune hypothyroidism [REF:imithot2024]. The trial represents the first attempt to use microbiome restoration as a therapeutic intervention for thyroid disease — moving beyond correlational studies to test whether fixing the gut can fix the thyroid.

Results are not yet available, and it would be premature to anticipate the outcome. The history of

FMT for conditions beyond *Clostridioides difficile* infection — as we saw with autism in Chapter 13 — has been a story of initial enthusiasm tempered by rigorous trial results. But the trial's existence is itself significant. It reflects a growing recognition that the gut-thyroid axis is not merely an academic curiosity but a potential therapeutic target.

15.4 Circadian Rhythms — Do Our Microbes Have Body Clocks?

The Clock Inside Every Cell

You have, inside virtually every cell in your body, a clock. Not a metaphorical clock — a literal molecular timekeeping mechanism, built from interlocking loops of gene expression, that cycles with a period of approximately 24 hours. This is the **circadian clock**, and it is one of the most ancient and conserved features of life on Earth. Versions of it exist in bacteria, fungi, plants, insects, and every animal ever studied. It evolved because the Earth rotates, and organisms that can anticipate the daily cycle of light and dark — rather than merely reacting to it — have a survival advantage. If you can gear up your metabolism before dawn, start producing digestive enzymes before your first meal, and prepare your immune defences before the nocturnal predators emerge, you are better equipped than an organism that has to start from scratch each time the environment changes.

In mammals, the master clock sits in a tiny cluster of about 20,000 neurons in the hypothalamus called the **suprachiasmatic nucleus (SCN)** — so named because it sits just above the optic chiasm, where the optic nerves cross. The SCN receives light information directly from the retina through specialised photoreceptive cells (not the rods and cones that provide vision, but a separate class of cells containing the photopigment melanopsin), and it uses this light signal to synchronise its internal oscillation to the external day-night cycle. From the SCN, timing signals radiate outward — through neural connections and hormonal signals, including the pineal gland's production of **melatonin** — to synchronise the peripheral clocks in every other tissue.

The molecular mechanism of the clock is an elegant negative feedback loop. A pair of proteins — **CLOCK** and **BMAL1** — bind together and activate the transcription of their own repressors, a pair of proteins called **PER** (period) and **CRY** (cryptochrome). PER and CRY accumulate, inhibit CLOCK/BMAL1, and shut down their own production. As PER and CRY are gradually degraded, the inhibition lifts, CLOCK/BMAL1 become active again, and the cycle restarts. One complete cycle takes approximately 24 hours — hence “circadian,” from the Latin *circa diem*, meaning “about a day.”

This clock does not merely tell time. It regulates time. Circadian genes control the timing of hormone secretion (cortisol peaks in the morning, melatonin peaks at night), immune function (inflammatory responses are stronger during the day), metabolism (insulin sensitivity is higher in the morning than at night), cell division, DNA repair, and gut function (gut motility, gastric acid secretion, and intestinal permeability all oscillate over 24 hours). When the clock is disrupted — by shift work, jet lag, chronic sleep deprivation, or constant light exposure — the health consequences are serious: increased risk of obesity, type 2 diabetes, cardiovascular disease, cancer, depression, and impaired immune function. The World Health Organisation has classified shift work involving circadian disruption as a probable carcinogen.

Given what we have learned about the microbiome's intimate involvement with metabolism, immunity, and hormonal signalling, you might expect the circadian clock and the gut microbiome to be connected. You would be right.

The Oscillating Microbiome

In 2014, Christoph Thaiss, Eran Elinav, and colleagues at the Weizmann Institute of Science in Israel published a study in *Cell* that demonstrated, for the first time, that the gut microbiome is not static over the course of a day [REF:thaiss2014]. It oscillates.

The team collected faecal samples from mice at multiple time points across the 24-hour cycle and found that the composition of the gut microbiome fluctuated rhythmically. The relative abundance of major bacterial phyla — Bacteroidetes and Firmicutes — oscillated in antiphase: when one was high, the other was low, and vice versa. Individual bacterial species showed their own rhythms, with different species peaking at different times of day. The functional capacity of the microbiome also oscillated — the metabolic pathways that were most active in the morning were different from those most active at night.

These oscillations were driven by the host's feeding rhythm. When mice ate on their normal schedule (mice are nocturnal, so they eat mainly at night), the microbial oscillations were robust and regular. When the feeding schedule was disrupted — by providing food only during the day, when mice normally sleep — the oscillations shifted accordingly. And when the host's circadian clock was genetically disrupted (using mice lacking the *PER1* and *PER2* genes, which are essential components of the molecular clock), the microbial oscillations were abolished. The gut bacteria were not keeping time on their own. They were being entrained by the host — specifically, by the rhythmic delivery of food and the rhythmic secretion of host-derived molecules (bile acids, antimicrobial peptides, mucus) that themselves follow the circadian clock.

The implications were provocative. If the microbiome oscillates, and if the microbiome produces metabolites that influence host physiology, then the *timing* of those metabolites matters. A short-chain fatty acid produced at 8 a.m. reaches a host whose metabolic and immune systems are in “morning mode.” The same molecule produced at 2 a.m. reaches a host in “night mode,” with different receptor sensitivities, different hormone levels, and different gene expression profiles. The effect of the metabolite could be entirely different depending on when it arrives.

Jet Lag, Dysbiosis, and Metabolic Disease

Thaiss's 2014 study did not stop at demonstrating that the microbiome oscillates. It showed what happens when the oscillations are disrupted — and the results were alarming.

The team induced jet lag in mice by shifting the light-dark cycle by eight hours — simulating a transatlantic flight. The jet-lagged mice showed disrupted microbial oscillations: the normal rhythmic cycling of bacterial composition was flattened or phase-shifted. This was not merely an interesting observation. The jet-lagged mice developed **glucose intolerance** and **increased body fat** — hallmarks of metabolic syndrome. And here is the critical experiment: when the researchers took the dysbiotic, arrhythmic microbiome from jet-lagged mice and transplanted it into germ-free mice, the recipients *also* developed glucose intolerance and obesity — even though the germ-free recipients had never experienced jet lag themselves. The metabolic dysfunction transferred with the microbiome.

The team then extended the finding to humans. They collected faecal samples from two human volunteers before, during, and after an international flight (Israel to the United States — an eight-to-ten-hour time zone shift). The post-flight microbiome showed the same disruption seen in jet-lagged mice: loss of normal community structure, expansion of bacterial taxa associated with

obesity and metabolic disease (notably members of the Firmicutes), and — when transplanted into germ-free mice — increased adiposity and glucose intolerance.

Two years later, Thaiss's group published a follow-up study in *Cell* that went deeper into the mechanism [REF:thaiss2016]. They showed that the diurnal oscillations of the gut microbiome drive global programming of the host's circadian transcriptome — the rhythmic turning on and off of genes in host tissues. In both the intestine and the liver, the microbiome's rhythmic metabolite production entrained oscillations in host gene expression, including genes involved in metabolism, drug processing, and immune function. When the microbial oscillations were disrupted (by antibiotics or germ-free conditions), the host's own transcriptional rhythms were disrupted too — and, remarkably, entirely new oscillations appeared in the host genome that were not present under normal conditions. The microbiome was not just riding the host's clock. It was helping to write the host's circadian programme.

The Metabolic Messenger: SCFAs and the Clock

How does the microbiome communicate with the circadian clock? The answer involves familiar molecules. Short-chain fatty acids — the same butyrate, acetate, and propionate that we have encountered throughout this book — appear to be key mediators.

In 2015, Vanessa Leone and colleagues at the University of Chicago demonstrated that the gut microbiome's influence on host circadian gene expression is mediated, in part, by SCFAs [REF:leone2015]. Working with germ-free and conventionally raised mice, they showed that germ-free mice had markedly impaired circadian clock gene expression in both the liver and the central clock — despite the presence of normal light-dark signals. The clock was ticking, but it was out of tune. When specific microbial metabolites, including butyrate, were administered to germ-free mice, circadian gene expression was partially restored.

The mechanism involves epigenetic regulation — changes in gene expression that do not alter the DNA sequence itself but change how accessible genes are for transcription. Butyrate, as we noted in Chapter 13, is a histone deacetylase inhibitor (HDACi). It modifies the protein scaffolding (histones) around which DNA is wound, loosening the structure and making genes more accessible for transcription. The circadian clock genes — CLOCK, BMAL1, PER, CRY — are themselves regulated by histone modifications. By producing butyrate in a rhythmic pattern (high after fibre-rich meals, low during fasting), the microbiome provides a rhythmic epigenetic signal that reinforces and fine-tunes the host's circadian gene expression.

Secondary bile acids contribute as well. Lithocholic acid — the secondary bile acid produced by microbial 7 α -dehydroxylation of chenodeoxycholic acid, as we discussed in section 15.1 — has been shown to influence circadian gene expression in intestinal cells. A 2024 study demonstrated that bacterially produced bile acid metabolites can directly modulate the period length of the circadian clock in host intestinal cells [REF:wang2024]. The microbiome's bile acid chemistry is not just a metabolic conversation with the liver. It is a timing signal.

Shift Work, Modern Life, and the Broken Clock

The circadian-microbiome connection has uncomfortable implications for modern life. Shift work, artificial lighting, irregular meal timing, chronic sleep deprivation, and transmeridian travel are all features of contemporary industrial society — and all of them disrupt circadian rhythms. If circadian disruption also disrupts the gut microbiome, and if a disrupted microbiome further impairs cir-

cadian function (as Thaïss's data suggest), then modern lifestyles may be creating a self-reinforcing cycle of chronobiological and microbial dysfunction.

The epidemiological data are consistent with this concern. Shift workers have higher rates of obesity, type 2 diabetes, cardiovascular disease, and certain cancers — and they also have altered gut microbiome composition. Whether the microbiome disruption is a cause, a consequence, or a mediator of the metabolic health effects of shift work is not yet clear. But the Thaïss jet-lag experiment — in which microbiome disruption alone, transferred to a healthy host, was sufficient to produce metabolic dysfunction — suggests that the microbiome is at least part of the causal chain.

There is also a simpler, more actionable implication: **when you eat may matter as much as what you eat**, at least for the microbiome. The gut microbiome's oscillations are driven by meal timing. Eating at irregular hours — late-night snacking, skipping breakfast, shifting meals by hours from day to day — disrupts the feeding-fasting cycle that entrains the microbial community. **Time-restricted eating** — confining food intake to a consistent window of roughly 8 to 12 hours per day — has been shown in animal models to restore microbial oscillations, improve metabolic parameters, and reduce the harmful effects of a high-fat diet, even without reducing total caloric intake. The benefits of time-restricted eating may be mediated, at least in part, through the restoration of normal microbiome rhythmicity.

Human evidence for time-restricted eating is growing but not yet definitive. Several randomised trials have shown modest benefits for weight loss, insulin sensitivity, and inflammatory markers, but the results are heterogeneous, and the contribution of the microbiome to these effects has not been isolated in most studies. As with so much in microbiome science, the animal evidence is compelling, the mechanistic logic is sound, and the human translation is lagging behind.

The Hormonal Ecosystem

This chapter has traced four conversations between the gut microbiome and the endocrine system — four ways in which the trillions of microbes in the gut participate in the hormonal regulation of the body.

Through bile acid metabolism, the microbiome tunes the FXR and TGR5 receptors that regulate fat storage, glucose metabolism, and energy expenditure — acting as a co-author of the metabolic programme that the liver executes. Through the estrobolome, it modulates circulating oestrogen levels, with implications for bone health, cardiovascular disease, and cancer risk. Through the gut-thyroid axis, it influences nutrient absorption, immune tolerance, and possibly the metabolism of thyroid hormones themselves — affecting the metabolic thermostat that sets the pace of every cell. And through its circadian oscillations, it helps synchronise the body's internal clocks, translating the rhythm of meals into the rhythm of gene expression in tissues throughout the body.

None of these conversations is fully understood. The bile acid story is the most advanced, with clear molecular mechanisms, validated receptors, and drugs in clinical trials. The estrobolome is biologically compelling but still resting largely on correlational evidence. The thyroid-microbiome connection is consistent but awaiting definitive interventional trials. And the circadian story, while backed by some of the most elegant experiments in the field, has yet to produce clear therapeutic applications in humans.

But the pattern is now unmistakable. The microbiome is not a passive inhabitant of the gut, digesting fibre and occupying space. It is a metabolically active, hormonally integrated component of the body's regulatory systems. It processes, modifies, and recycles the body's own hormones. It produces signalling molecules that activate hormone receptors. It sets the conditions under which hormonal messages are sent, received, and interpreted. In every meaningful sense, the gut microbiome is an endocrine organ — unstructured, distributed, and microbial in nature, but performing endocrine functions as surely as the thyroid, the adrenal glands, or the ovaries.

This is the last chapter of Part III — the section of this book devoted to how the microbiome talks to the body. We began, in Chapter 12, with the immune negotiation: the lifelong conversation between microbe and host that calibrates the body's defences. We continued, in Chapters 13 and 14, with the gut-brain axis: the neural, chemical, and behavioural dialogue between the gut and the mind. And we end here, with the endocrine microbiome: the hormonal conversation that links the gut to the liver, the reproductive system, the thyroid, and the body's internal clocks.

In Part IV, we change direction. We leave the world of bacteria and enter the world of viruses — not the viruses that cause disease, but the viruses that live *within* the microbiome. The bacteriophages. They are the most abundant biological entities on Earth, they are present in your gut in staggering numbers, and they shape the bacterial community in ways we are only beginning to understand. They are the missing players in the story we have been telling — and they have been there all along.

Chapter 15 — Key References

- [REF:sayin2013] Sayin, S.I., Wahlström, A., Felin, J., Jäntti, S., Marschall, H.U., Bamberg, K., Angelin, B., Hyötyläinen, T., Orešič, M. & Bäckhed, F. (2013). Gut microbiota regulates bile acid metabolism by reducing the levels of tauro-beta-muricholic acid, a naturally occurring FXR antagonist. *Cell Metabolism*, 17(2), 225–235. DOI: 10.1016/j.cmet.2013.01.003
- [REF:wahlstrom2016] Wahlström, A., Sayin, S.I., Marschall, H.U. & Bäckhed, F. (2016). Intestinal crosstalk between bile acids and microbiota and its impact on host metabolism. *Cell Metabolism*, 24(1), 41–50. DOI: 10.1016/j.cmet.2016.05.005
- [REF:plottel2011] Plottel, C.S. & Blaser, M.J. (2011). Microbiome and malignancy. *Cell Host & Microbe*, 10(4), 324–335. DOI: 10.1016/j.chom.2011.10.003
- [REF:flores2012] Flores, R., Shi, J., Fuhrman, B., Xu, X., Veenstra, T.D., Gail, M.H., Gajer, P., Ravel, J. & Goedert, J.J. (2012). Fecal microbial determinants of fecal and systemic estrogens and estrogen metabolites: a cross-sectional study. *Journal of Translational Medicine*, 10, 253. DOI: 10.1186/1479-5876-10-253
- [REF:larnder2025] Larnder, A.H., Bhatt, A.P. & Engstrand, L. (2025). The estrobolome: estrogen-metabolizing pathways of the gut microbiome and their relation to breast cancer. *International Journal of Cancer*. DOI: 10.1002/ijc.35427
- [REF:li2025] Li, X. et al. (2025). The testobolome in microbial testosterone metabolism and human health. *npj Biofilms and Microbiomes*. DOI: 10.1038/s41522-025-00861-0
- [REF:knezevic2020] Knezevic, J., Starchl, C., Tmava Berisha, A. & Amrein, K. (2020). Thyroid-gut-axis: how does the microbiota influence thyroid function? *Nutrients*, 12(6), 1769. DOI: 10.3390/nu12061769

- [REF:sun2024] Sun, J. et al. (2024). Recent advances in gut microbiota and thyroid disease: pathogenesis and therapeutics in autoimmune, neoplastic, and nodular conditions. *Frontiers in Cellular and Infection Microbiology*, 14, 1465928. DOI: 10.3389/fcimb.2024.1465928
- [REF:shi2024] Shi, T.T. et al. (2024). Intestinal microbiota regulates the gut-thyroid axis: the new dawn of improving Hashimoto thyroiditis. *Clinical and Experimental Medicine*, 24, 39. DOI: 10.1007/s10238-024-01292-x
- [REF:chen2025] Chen, Y. et al. (2025). Gut microbiota in hypothyroidism: pathogenic mechanisms and opportunities for precision microbiome interventions. *Frontiers in Microbiology*, 16, 1661211. DOI: 10.3389/fmicb.2025.1661211
- [REF:imithot2024] de Groot, P. et al. (2024). The IMITHOT trial: fecal microbiota transplantation in subclinical autoimmune hypothyroidism. *ClinicalTrials.gov* / study protocol, Amsterdam UMC.
- [REF:thaiss2014] Thaiss, C.A., Zeevi, D., Levy, M., Zilberman-Schapira, G., Suez, J., Tengeler, A.C., Abramson, L., Katz, M.N., Korem, T., Zmora, N., Kuperman, Y., Biton, I., Gilad, S., Harmelin, A., Shapiro, H., Halpern, Z., Segal, E. & Elinav, E. (2014). Transkingdom control of microbiota diurnal oscillations promotes metabolic homeostasis. *Cell*, 159(3), 514–529. DOI: 10.1016/j.cell.2014.09.048
- [REF:thaiss2016] Thaiss, C.A., Levy, M., Korem, T., Dohnalová, L., Shapiro, H., Jaitin, D.A., David, E., Winter, D.R., Gury-BenAri, M., Tatirovsky, E., Tuganbaev, T., Federici, S., Zmora, N., Zeevi, D., Dori-Bachash, M., Pevsner-Fischer, M., Kartvelishvili, E., Brandis, A., Harmelin, A., Shibolet, O., Halpern, Z., Honda, K., Amit, I., Segal, E. & Elinav, E. (2016). Microbiota diurnal rhythmicity programs host transcriptome oscillations. *Cell*, 167(6), 1495–1510.e12. DOI: 10.1016/j.cell.2016.11.003
- [REF:leone2015] Leone, V., Gibbons, S.M., Martinez, K., Hutchison, A.L., Huang, E.Y., Cham, C.M., Pierre, J.F., Heneghan, A.F., Nadimpalli, A., Hubert, N., Zale, E., Wang, Y., Huang, Y., Theriault, B., Dinner, A.R., Musch, M.W., Kudsk, K.A., Prendergast, B.J., Gilbert, J.A. & Chang, E.B. (2015). Effects of diurnal variation of gut microbes and high-fat feeding on host circadian clock function and metabolism. *Cell Host & Microbe*, 17(5), 681–689. DOI: 10.1016/j.chom.2015.03.006
- [REF:wang2024] Wang, Y. et al. (2024). Gut microbiome–produced bile acid metabolite lengthens the circadian period in host intestinal cells. *Proceedings of the National Academy of Sciences*. DOI: 10.1073/pnas.2506313123

Additional references to be verified and expanded during manuscript preparation.

Chapter 16: Bacteriophages — The Viruses Within the Microbiome

For fifteen chapters, we have been telling the story of the human microbiome as if it were principally a story about bacteria — and, in Chapter 10, about fungi. We have mapped the bacterial communities of the gut, the skin, the lungs, the mouth, and the urogenital tract. We have followed the chemical conversations between these microbes and the immune system, the brain, and the endocrine system. We have built, piece by piece, a picture of the human body as an ecosystem — a partnership between human cells and microbial ones, negotiated over evolutionary time and maintained through continuous molecular dialogue.

But we have been telling an incomplete story. Because within that microbial ecosystem, there is another layer of life — if “life” is even the right word — that we have barely mentioned. The bacteria in your gut are not undisturbed. They are under constant siege. They are being infected, hijacked, and killed by organisms even smaller than themselves: viruses that prey exclusively on bacteria. They are called **bacteriophages** — from the Greek *phagein*, to devour — and they are, by any reasonable measure, the most abundant biological entities on the planet.

There are an estimated 10^{31} bacteriophages on Earth. Written out, that is a one followed by thirty-one zeroes: 10,000,000,000,000,000,000,000,000,000. Ten million trillion trillion. If you laid every phage on the planet end to end, the line would stretch for about 200 million light-years — far enough to span several galaxy clusters. There are more phages on Earth than there are stars in the observable universe. More phages than every other biological entity — every bacterium, every fungus, every plant and animal cell — combined [REF:mushegian2020].

And a substantial number of them live inside you.

Your gut alone harbours an estimated 10^9 to 10^{10} virus-like particles per gram of intestinal contents — a number roughly equal to, or exceeding, the bacterial population in the same gram [REF:shkoporov2019]. In a typical human colon, that translates to something on the order of 10^{12} to 10^{13} individual phage particles — trillions of viruses, quietly going about the business of infecting, replicating within, and destroying the bacteria that we have spent the previous chapters getting to know. They are the unseen regulators of the microbial ecosystem, the invisible predators that shape which bacteria thrive and which are held in check, the molecular shuttles that move genes between bacterial species with a promiscuity that would make a geneticist nervous. And until remarkably recently, almost nobody in biomedical science was paying attention to them.

This chapter and the next are about putting that right. Here, we meet the phages: what they are, how they were discovered and then forgotten, what the human “phageome” looks like, how phages regulate bacterial populations, and how they blur the boundary between virus and host in ways that reshape the entire microbiome. In Chapter 17, we will explore the even more surprising discovery that phages do not merely interact with bacteria — they interact directly with us, crossing epithelial barriers, engaging the immune system, and participating in the body’s defences in ways that are still being mapped.

Together, these two chapters complete the picture. The human microbiome is not a community of

bacteria with a smattering of fungi. It is a community of bacteria, fungi, and the viruses that infect them — a tripartite ecosystem in which the phages are not peripheral players but central architects of community structure. To understand the microbiome without understanding phages is like trying to understand a forest without knowing about the insects that pollinate the trees, consume the leaves, and decompose the wood. The forest might look the same at a glance. But the forces that shape it would be invisible.

16.1 What Are Phages and Why Were They Forgotten?

Smaller Than Small

A bacteriophage is a virus that infects bacteria. That sentence is simple enough, but each word in it carries weight, so let us unpack it.

A virus is not, strictly speaking, alive. It has no metabolism — it cannot generate energy, synthesise proteins, or reproduce on its own. It is, in essence, a set of genetic instructions wrapped in a protein coat, sometimes with an additional lipid envelope. A virus can do exactly one thing: enter a host cell, commandeer the host's molecular machinery, and use it to produce copies of itself. Outside a host cell, a virus is inert — a particle, not an organism. The debate about whether viruses are “alive” has raged since they were first described, and it remains unresolved. But for our purposes, the functional definition is what matters: a virus is a genetic parasite that can only replicate inside a living cell.

A bacteriophage is a virus whose host cell is a bacterium. Where the viruses that concern human medicine — influenza, SARS-CoV-2, HIV — infect human cells, phages infect bacterial cells. They are extraordinarily specific: a given phage typically infects only a single bacterial species, often only particular strains within that species. A phage that infects *Escherichia coli* will not infect *Bacteroides*, and vice versa. This specificity is determined by the molecular lock-and-key interaction between the phage's tail fibres (the structures it uses to attach to a host cell) and specific receptor molecules on the bacterial surface. Change the receptor, and the phage cannot attach. Change the lock, and the key does not fit.

Phages come in a remarkable variety of shapes and sizes. The most iconic are the **tailed phages** — formerly classified in the order Caudovirales, now reclassified into multiple families under the class Caudoviricetes following a major taxonomic revision by the International Committee on Taxonomy of Viruses. These phages look, under an electron microscope, like nothing else in biology: an icosahedral head (a geometric structure with twenty triangular faces, resembling a faceted jewel) sitting atop a cylindrical tail, often with spidery tail fibres extending from the base. They look, frankly, like miniature lunar landers — and they function in a conceptually similar way. The head contains the phage's genome. The tail serves as the injection apparatus. The tail fibres locate and lock onto the bacterial surface. And then the phage does something remarkable: it contracts its tail sheath, driving a hollow needle through the bacterial cell wall and injecting its DNA into the cell, like a syringe pushing through skin. The protein coat remains outside. Only the genetic material enters.

Other phage morphologies exist. The **Microviridae** are small, tailless, icosahedral phages with single-stranded DNA genomes — simpler, more compact, and abundant in the gut virome. Filamentous phages are long, thin, flexible rods. But the tailed phages dominate in diversity and are the most thoroughly studied, so they will serve as our primary reference throughout this chapter.

The genomes of phages are typically small — far smaller than bacterial genomes. A typical tailed phage carries between 30,000 and 200,000 base pairs of DNA, compared to the several million base pairs in a typical bacterial genome. Yet those compact genomes encode everything the phage needs to hijack a bacterium: genes for attaching to the host, injecting DNA, shutting down host gene expression, replicating the phage genome, assembling new phage particles, and — in the final act — lysing the bacterial cell, bursting it open to release the newly made phages into the environment, each one ready to find a new host and begin the cycle again. A single infected bacterium can release anywhere from a few dozen to several hundred new phage particles, depending on the phage and the host. The bacterial cell does not survive the experience.

This is the **lytic cycle** — the classic phage life cycle described in every virology textbook. Attach, inject, replicate, lyse, repeat. It is fast (many phages complete the cycle in under an hour), efficient (each infected cell produces many copies), and lethal for the host bacterium. It is also, as we shall see, only half the story.

The Discovery That Was Discovered Twice

The story of how bacteriophages were found — and then lost — is one of the more frustrating episodes in the history of biology.

In 1915, the English bacteriologist Frederick Twort, working at the Brown Institution in London, observed something puzzling. He was attempting to grow vaccinia virus (the virus used in smallpox vaccination) on agar plates seeded with bacteria, when he noticed that some of his bacterial colonies developed glassy, transparent areas — zones where the bacteria appeared to have dissolved. The agent responsible could pass through filters fine enough to exclude bacteria, could be transmitted from one colony to another, and could be diluted through multiple passages, suggesting that it was replicating. Twort published his findings in *The Lancet*, noting cautiously that the phenomenon might be caused by “an ultra-microscopic virus” — but he was uncertain, and he did not pursue the work further. The First World War intervened, funding dried up, and Twort moved on to other projects [REF:twort1915].

Two years later, in 1917, Félix d’Hérelle — a self-taught French-Canadian microbiologist working at the Pasteur Institute in Paris — made the same discovery independently, but with considerably more flair and considerably less caution. D’Hérelle was investigating an outbreak of dysentery among French soldiers and noticed that filtrates from the stool of recovering patients could lyse cultures of *Shigella*, the bacterium responsible for the disease. The lytic agent could be serially passaged: a tiny drop of the filtrate added to a fresh bacterial culture would clear it overnight, and a drop from that cleared culture would clear the next, and so on. Each passage produced more of the agent, confirming that it was replicating — and that it required bacteria to do so [REF:dherelle1917].

D’Hérelle grasped immediately what he had found. He named the agent “bacteriophage” — eater of bacteria — and published his results with a bold interpretive framework: these were viruses that parasitised bacteria, they were ubiquitous in nature, and they could be harnessed to treat bacterial infections. Within five years, he was using phage preparations to treat dysentery, cholera, and plague in humans, reporting dramatic successes. The earliest documented therapeutic use of phages in the United States came in 1922 [REF:aswani2021].

The scientific community was electrified — but also deeply divided. Jules Bordet, the Nobel laureate immunologist (also at the Pasteur Institute), disputed d’Hérelle’s interpretation. Bordet argued that the lytic agent was not a living virus but a bacterial enzyme — an autolytic substance produced

by the bacteria themselves. The Bordet-d'Hérelle debate consumed microbiological meetings for a decade and was conducted with a personal acrimony that went well beyond scientific disagreement. D'Hérelle, who lacked formal credentials and made grandiose claims, was an easy target. Bordet, who had the Nobel Prize and the institutional authority, was a formidable opponent.

The debate was ultimately resolved — but not quickly enough to save phage research from what happened next.

The Eclipse

In 1928, Alexander Fleming noticed a mould killing bacteria on a petri dish. By the 1940s, penicillin was in mass production, and the antibiotic era had begun.

The effect on phage research was devastating. Antibiotics were broad-spectrum — a single drug could kill many different bacterial species. They were chemical, standardisable, and could be manufactured in industrial quantities. They were patentable. Phage therapy, by contrast, was narrow-spectrum (each phage killed only one bacterial species or strain), required fresh preparation, was poorly standardised, and the results reported from clinical trials were inconsistent — partly because the science of phage biology was still immature, and partly because some early therapeutic claims had been exaggerated.

In the West, the choice was made. Antibiotics won. Phage therapy was abandoned by mainstream medicine in Europe and North America by the early 1950s. The phages themselves were not forgotten entirely — they became indispensable tools for the new science of molecular biology. Max Delbrück, a German physicist who had emigrated to the United States, assembled a group of brilliant researchers — the informal “phage group” that included Salvador Luria and Alfred Hershey — who used phages as model systems to crack open the fundamental mechanisms of genetics. It was Hershey and Martha Chase who, in 1952, used phages to prove definitively that DNA, not protein, carried genetic information — one of the pivotal experiments in the history of biology. Delbrück, Luria, and Hershey shared the Nobel Prize in Physiology or Medicine in 1969 for their work on “the replication mechanism and the genetic structure of viruses” [REF:norrby2008].

The irony is considerable. Phages were central to the birth of molecular biology — and yet phage biology itself, as an ecological and medical discipline, withered. In the Soviet Union and its satellite states, phage therapy continued — the Eliava Institute in Tbilisi, Georgia, maintained an active phage therapy programme throughout the Cold War, and phage preparations remain available over the counter in Russian and Georgian pharmacies to this day. But in the West, phages became laboratory curiosities, useful for experiments but irrelevant to medicine. The idea that phages might be important components of the human body's own microbial ecosystem — that they might be doing something in our guts, on our skin, in our lungs — was not so much rejected as simply never considered.

That began to change only in the 2000s, when metagenomic sequencing — the same technology that revealed the true diversity of the bacterial microbiome (Chapter 4) — was turned for the first time on the viral fraction of the gut contents. What it revealed was staggering.

16.2 The Human “Phageome” — Scale, Diversity, and Dynamics

Lifting the Veil

Studying viruses in the gut is harder than studying bacteria, for reasons that are partly technical and partly conceptual.

The technical problem is extraction. When researchers want to study the bacterial microbiome, they can collect a stool sample, extract all the DNA, and sequence it. Most of that DNA will be bacterial, and the well-established databases of bacterial genomes allow researchers to identify the species present with reasonable confidence. But viruses — and phages in particular — are far smaller than bacteria, their genomes are far smaller, and they are present in a complex mixture with an overwhelming excess of bacterial and human DNA. To study the phageome, researchers must first separate the viral particles from everything else. The standard approach involves filtering the sample through membranes with pores small enough to exclude bacteria (typically 0.2 or 0.45 micrometres), treating the filtrate with nucleases to destroy any free-floating DNA that is not protected inside a viral capsid, and then extracting and sequencing the DNA from the remaining virus-like particles (VLPs). The process is laborious, subject to contamination at every step, and recovers only a fraction of the viruses actually present — particularly missing those that are integrated into bacterial genomes as prophages (a category we will return to in section 16.4).

The conceptual problem is worse. When the first gut virome studies were conducted in the mid-2000s, researchers found that the vast majority of the viral sequences they recovered — typically 75 to 90 per cent — matched nothing in existing databases. They were completely unknown. In bacterial metagenomics, you can usually assign most sequences to known phyla, families, or genera, because decades of culture-based microbiology and genome sequencing have built up a substantial reference library. For viruses, no such library existed. The gut virome was, and to a significant extent still is, what researchers call **viral dark matter** — sequences that are clearly viral in origin (they have the hallmarks of viral genomes) but cannot be assigned to any known viral group because they have never been seen before [REF:pargin2023].

This is worth pausing on, because it reveals something profound about the state of our knowledge. We are not talking about a few unknown species. We are talking about the majority of all viruses in the human gut having no match in any database on Earth. It is as if you conducted a census of a major city and found that 80 per cent of the residents did not appear in any government record, any immigration database, any birth certificate, anywhere. They were there — you could count them, you could see their houses — but you had no idea who they were or where they had come from.

Despite these challenges, the past fifteen years have produced an increasingly detailed — if still incomplete — portrait of the human phageome. And the portrait is remarkable.

The Numbers

The human gut virome is dominated by bacteriophages. Other viruses are present — including eukaryotic viruses that infect human cells, and plant viruses ingested with food — but phages account for the vast majority of the viral particles recovered from faecal samples and intestinal biopsies.

Estimates of phage abundance in the gut have varied with methodology. Early studies using electron microscopy to count virus-like particles in faecal samples reported approximately 10^8 to 10^9 VLPs per gram. More recent approaches using quantitative sequencing with spike-in controls — adding a known quantity of a reference virus to the sample before extraction, so that the abso-

lute abundance of other viruses can be calculated by comparison — have yielded higher estimates, on the order of 10^{10} VLPs per gram, roughly five-fold higher than the microscopy-based figures [REF:shkoporov2019]. The discrepancy likely reflects the limitations of microscopy (which can miss small or fragile particles) rather than a genuine increase in phage numbers. By any estimate, though, the order of magnitude is clear: the number of phages in the gut is comparable to, or exceeds, the number of bacteria.

This is a staggering realisation when set alongside the previous chapters. Every bacterial community we have described — in the gut, on the skin, in the lungs — exists under constant viral pressure. The bacteria are not merely competing with each other for nutrients and space. They are simultaneously defending themselves against a vast and diverse army of phages, each one exquisitely adapted to exploit a particular bacterial host. The microbiome is not a garden. It is a garden under perpetual siege.

crAssphage — The Virus Hiding in Plain Sight

In 2014, a team led by Bas Dutilh at Radboud University in the Netherlands made one of the most extraordinary discoveries in virome research — not by collecting new samples or developing new technology, but by reanalysing data that was already sitting in public databases [REF:dutilh2014].

Dutilh's approach was computational. His team took existing human faecal metagenomes — datasets of DNA sequences that had already been generated by other research groups — and used a method called cross-assembly to identify sequences that appeared repeatedly across multiple samples but did not match anything in reference databases. Cross-assembly works by pooling reads from many samples and assembling them together, allowing sequences that are too fragmented in any single sample to be reconstructed into longer, more interpretable contigs. The method's name — crAss, for **cross-Assembly** — would give the discovery its own name.

What emerged from the analysis was a single, circular viral genome of approximately 97,000 base pairs — large for a phage, encoding 80 predicted proteins, the vast majority of which had no detectable homology to anything in existing databases. The genome was clearly viral: it had the structural hallmarks of a phage, including genes for capsid proteins and DNA packaging. But it was unlike any known virus.

And it was everywhere. When Dutilh's team searched for this genome in published human faecal metagenomes, they found it in the majority of samples — present in roughly 50 to 70 per cent of all individuals in Western populations. In some samples, it accounted for up to 90 per cent of all viral sequences. It was, by a wide margin, the single most abundant virus ever identified in the human gut — estimated to be roughly six times more abundant in public metagenomes than all other known phages combined. It had been lurking in datasets around the world, in plain sight, for years. Nobody had noticed it because nobody recognised it. Its genes were too novel to match anything in the reference databases, so standard bioinformatic pipelines had simply classified it as “unknown” and moved on.

Dutilh named it **crAssphage** — after the cross-assembly method that had revealed it.

Subsequent work identified crAssphage's host. Using a combination of bioinformatic predictions and experimental confirmation, researchers established that crAssphage infects members of the genus *Bacteroides* — one of the most abundant bacterial groups in the human gut, the same genus we encountered in Chapter 5 as a dominant member of the colonic community. Specifically, it

appears to target species including *Phocaeicola vulgatus* (formerly *Bacteroides vulgatus*) and *Bacteroides stercoris*. In 2025, Schmidtke and colleagues published a landmark study in *Cell Host & Microbe* that revealed something unexpected about crAssphage's biology: it is not a conventional phage at all, but a **phage-plasmid** — a hybrid genetic element that can exist as a linear extra-chromosomal molecule inside the bacterial cell, replicating stably without killing its host, and only occasionally switching to a lytic mode that produces new phage particles [REF:schmidtke2025]. This discovery helped explain one of the puzzles of crAssphage biology: how a phage could be so enormously abundant without visibly decimating its host population. The answer, it seems, is that crAssphage has found a way to have it both ways — persisting quietly inside the bacterial cell most of the time, and producing new particles only intermittently.

CrAssphage has since been recognised as the founding member of a much larger group of related phages — the **Crassvirales** — that together constitute one of the most abundant viral groups in the human gut, particularly in industrialised populations. Its discovery was a watershed moment for virome research, not only because of the biological significance of the virus itself, but because of what it revealed about the limits of our knowledge. The most abundant virus in the human body had been completely unknown until 2014. If something that common could be missed, what else was hiding in the dark matter?

The Gut Virome Database

The problem of viral dark matter spurred efforts to build comprehensive reference databases for the human gut virome. In 2020, Ann Gregory, Matthew Sullivan, and colleagues published the **Gut Virome Database (GVD)** — a curated collection of 33,242 unique viral populations identified from 2,697 viral and microbial metagenomes drawn from 1,986 individuals across 16 countries [REF:gregory2020]. The GVD improved the rate at which viral sequences from new gut samples could be identified: compared with the standard RefSeq viral database maintained by NCBI, the GVD improved viral detection rates nearly 182-fold.

The GVD also revealed important patterns. The gut virome shows clear **age-dependent dynamics**: infants have viromes dominated by different phage groups than adults, and virome diversity changes across the human lifespan. The virome is **highly personalised** — each individual harbours a distinct community of phages, as individual-specific as a fingerprint. And the virome is **remarkably stable** over time in healthy adults. Minot and colleagues, in a pioneering 2011 study that sequenced 56 billion bases of viral DNA from 24 faecal samples collected from four individuals over 2.5 years, found that approximately 80 per cent of viral types persisted throughout the entire study period [REF:minot2011]. Your virome, like your bacterial microbiome, is your own — and it stays with you.

But the virome is not static. Minot's study also showed that the virome responds to dietary changes: individuals placed on the same controlled diet showed convergence in their virome composition, suggesting that diet shapes the virome just as it shapes the bacterial community, presumably because the virome tracks the bacteria that diet selects for. When your dietary choices reshape your bacterial community (as we discussed in Chapter 14), they reshape the viral community too — because the phages follow their hosts.

The Major Players

Beyond crAssphage and the Crassvirales, the human gut phageome is composed of several major groups. The **Caudoviricetes** (the tailed phages, the most morphologically diverse group of phages

known) are well represented, including members of the former families Siphoviridae, Myoviridae, and Podoviridae — designations that, while now taxonomically obsolete, remain in common use in the virome literature because the new nomenclature is still being adopted. The **Microviridae** — small, tailless, single-stranded DNA phages — are also abundant and show high mutation rates, suggesting rapid evolution within the gut environment. Together with the Crassvirales, these groups account for the majority of identifiable phages in the human gut, though the caveat about viral dark matter means that a substantial fraction of the virome remains uncharacterised.

What is clear, even from this incomplete picture, is that the gut phageome is not a random collection of viruses. It is a structured community, shaped by the bacterial community it parasitises, stable over time in healthy individuals, responsive to environmental change, and — as we are about to see — actively involved in regulating the very bacterial populations it depends on.

16.3 Phages as Regulators of Bacterial Populations

The Predator and the Prey

In ecology, the relationship between predator and prey is one of the fundamental forces that maintains biodiversity. On the African savanna, lions do not merely kill wildebeest — they shape the wildebeest population, preventing it from growing so large that it overgrazes the grasslands and outcompetes other herbivores. The lions, by suppressing the dominant herbivore, create space for zebras, gazelles, and smaller grazers that would otherwise be crowded out. Predation, paradoxically, maintains diversity.

The same principle operates in the gut — except the predators are phages and the prey are bacteria.

The ecological model that describes this dynamic is called **kill-the-winner**. The name is self-explanatory. In a diverse bacterial community, individual species compete for nutrients and space. When one species begins to dominate — winning the competition, growing faster than its neighbours — it becomes a larger target for phages. Phages that specialise in infecting that dominant species encounter more hosts, replicate faster, and increase in number, exerting intensifying predatory pressure on the winning species. As the dominant population is cut back, the competitive pressure on other species is relieved, and they can expand into the vacated niche. The result is a dynamic equilibrium — an oscillation in which no single bacterial species can achieve lasting dominance because success itself attracts the attention of predators.

This is, in mathematical terms, a variant of the **Lotka-Volterra equations** — the classic predator-prey model first described in the 1920s by Alfred Lotka and Vito Volterra to explain oscillations in fish populations in the Adriatic Sea. In the phage-bacteria version, bacterial population growth is checked by phage predation, phage numbers increase as bacterial hosts become available, phage numbers decline as they deplete their hosts, and the cycle repeats. The oscillations may be rapid — bacteria can double in as little as twenty minutes, and a phage lytic cycle can complete in under an hour — producing a dynamic, constantly shifting balance that, averaged over time, maintains a diverse bacterial community.

Allen and colleagues demonstrated a version of this dynamic in 2011, studying the swine gut microbiome. They observed an inverse relationship between the abundance of *Streptococcus* bacteria and the abundance of *Streptococcus*-infecting phages: when the phages were abundant, the bacteria were suppressed, and when the bacterial population recovered, phage numbers surged in response

— the classic oscillation predicted by kill-the-winner dynamics [REF:allen2011].

In the human gut, direct evidence for kill-the-winner has been harder to obtain, because the system is vastly more complex — dozens of bacterial species, hundreds of phage types, all interacting simultaneously. But the circumstantial evidence is compelling. Gut virome studies consistently show that phage abundance correlates with the abundance of their bacterial hosts, and that shifts in bacterial community composition are mirrored by corresponding shifts in the virome. The phages are tracking the bacteria, and the bacteria are responding to the phages. The two communities are locked in a continuous molecular dance.

Maintaining Diversity

The kill-the-winner dynamic has a profound ecological consequence: it promotes bacterial diversity. By preventing any single species from dominating the gut ecosystem, phages create a more varied community — one in which many species coexist, each held below the population threshold at which its specialist phages would erupt in force.

This matters for the host. We have seen throughout this book that microbial diversity is, broadly speaking, a good thing for human health. A diverse gut microbiome produces a wider range of metabolites, provides more robust colonisation resistance against pathogens (Chapter 5), offers more comprehensive immune education (Chapter 12), and is more resilient to perturbation. Reduced diversity is a hallmark of dysbiosis — the disrupted microbial state associated with inflammatory bowel disease, obesity, and a host of other conditions. If phages help maintain that diversity, then they are not merely parasites of the bacterial community. They are, from the host's perspective, regulators — gardeners of the ecosystem, selectively pruning the most aggressive species and maintaining the conditions under which a diverse community can flourish.

Think of it this way. Imagine a garden in which many species of plant are growing together. One species — call it the weed — grows faster than all the others and, if left unchecked, would eventually take over the entire plot, crowding out the flowers, the herbs, and the vegetables. Now introduce a pest that specifically targets that weed, eating its leaves and slowing its growth. The weed is no longer able to dominate. The other plants have room to grow. The garden, counterintuitively, is healthier *because* of the pest. The phages in the gut play the role of that pest — and the garden is your microbiome.

When the Balance Breaks

If phages maintain bacterial diversity, then disruptions to the phageome might contribute to the disrupted bacterial communities seen in disease. This hypothesis is increasingly supported by evidence. Several studies have found that patients with inflammatory bowel disease (IBD) have altered gut viromes — typically showing increased phage diversity and increased lytic activity compared with healthy controls. The interpretation is that, in the inflamed gut, prophages that had been lying dormant inside bacterial cells are induced to enter the lytic cycle (a process we will examine in the next section), leading to waves of bacterial lysis that further destabilise an already dysbiotic community. The phages, in this scenario, shift from being stabilising regulators to being amplifiers of disruption — a distinction that depends entirely on the ecological context.

This is a recurring theme in microbiome science, and one we have encountered before: the same biological agent can be beneficial in one context and harmful in another. A phage that maintains

diversity by pruning a dominant species in a healthy gut can exacerbate dysbiosis by lysing essential commensals in an inflamed one. The agent has not changed. The ecosystem around it has.

16.4 Lysogeny: When Viruses Become Part of the Bacterial Genome

The Other Life Cycle

Everything we have discussed so far — the lytic cycle, the kill-the-winner dynamic, the regulation of bacterial populations through predation — describes what happens when a phage behaves as a straightforward parasite: infecting a bacterium, replicating inside it, destroying it, and moving on. This is the lytic lifestyle, and it is dramatic, fast, and lethal.

But many phages — perhaps the majority of phages in the human gut — are capable of something altogether more subtle. Instead of immediately destroying their host, they can integrate their genome into the bacterial chromosome and go silent. The phage DNA is stitched into the host's DNA, becoming a physical part of the bacterial genome. When the bacterium divides, it copies the phage DNA along with its own, passing it to both daughter cells. The phage genome — now called a **pro-phage** — is replicated for free, inherited by every descendant, and can persist in this dormant state for hundreds or thousands of bacterial generations, adding nothing to the cell's burden and causing no harm.

This is **lysogeny**, and a phage capable of it is called a **temperate phage**. A bacterium carrying a prophage is called a **lysogen**.

The analogy that microbiologists often reach for is a sleeper agent — a spy who has infiltrated an organisation, assumed a normal identity, and lives quietly among its members, indistinguishable from a genuine employee, for years. The agent is still a spy. The potential for betrayal is still there. But as long as conditions are stable, the agent does nothing — simply existing within the host, biding time.

The numbers are striking. Analysis of sequenced bacterial genomes suggests that prophages are extraordinarily common. The majority of gut bacteria carry at least one integrated prophage, and many carry several. Some estimates suggest that prophage sequences account for 10 to 20 per cent of total bacterial genomic content in certain species. The gut is not merely a battlefield between phages and bacteria — it is a landscape in which a large fraction of the bacterial population has been quietly infiltrated by phage DNA that has become, for the moment, part of the bacterial genome itself.

The Switch

The prophage's silence is conditional. Under certain circumstances, the dormant phage genome reactivates — a process called **induction**. The phage DNA excises itself from the bacterial chromosome, the lytic cycle begins, the bacterium is destroyed, and a burst of new phage particles is released. The sleeper agent has woken up.

What triggers induction? The most thoroughly studied mechanism involves the bacterial **SOS response** — a DNA damage repair system that bacteria activate when their DNA is threatened. When a bacterium encounters an agent that damages its DNA — ultraviolet light, reactive oxygen species, certain chemicals — it activates the SOS response, which includes the production of a protease

called RecA. RecA, in the course of its DNA repair duties, cleaves a regulatory protein called the CI repressor, which is the molecular switch that keeps the prophage silent. With CI destroyed, the prophage is de-repressed, and the lytic programme activates.

In the gut, the most clinically relevant trigger for prophage induction is **antibiotics** — a connection that will become central in Chapter 18. Many antibiotics work by damaging bacterial DNA or interfering with DNA replication, and they activate the SOS response as a side effect. Allen and colleagues, in the same 2011 study that demonstrated kill-the-winner dynamics in the swine gut, found that antibiotic treatment (a cocktail of chlortetracycline, sulfamethazine, and penicillin) significantly increased the abundance of phage-related sequences in the gut virome — consistent with widespread prophage induction [REF:allen2011]. The antibiotics were not only killing bacteria directly; they were waking up the dormant phages inside the surviving bacteria, triggering a secondary wave of bacterial destruction. The medication was, in effect, arming the enemy within.

Other induction triggers relevant to the gut include **inflammatory signals** — reactive oxygen species produced by the immune system during intestinal inflammation can activate the SOS response — and **dietary changes** that alter the chemical environment of the gut. The implication is that prophage induction is not a rare event triggered only by extreme stress. It is a regular occurrence in the dynamic, chemically complex environment of the gut, contributing to the ongoing turnover and reshaping of the bacterial community.

Piggyback the Winner

If lysogeny seems disadvantageous for the host bacterium — a ticking time bomb embedded in its genome — then why is it so common? Why haven't bacteria evolved to reject prophages entirely?

Part of the answer is that they cannot always do so. Phage integration is, from the bacterium's perspective, an infection — an imposed event, not a choice. But the persistence of lysogeny across evolutionary time suggests that it is not purely parasitic. In many cases, carrying a prophage confers genuine advantages on the host bacterium.

The most straightforward benefit is **superinfection immunity**. A bacterium carrying a particular prophage is immune to subsequent infection by the same phage or closely related phages. The CI repressor protein that keeps the prophage dormant also represses the DNA of any new phage of the same type that tries to enter the cell. The lysogen is, in effect, vaccinated — protected from the very virus that has infiltrated it. In an environment saturated with phages, this immunity is a significant survival advantage. A lysogen can thrive in conditions where its non-lysogenised neighbours are being killed by lytic phages.

But the benefits go further. Prophage genomes often carry genes that have nothing to do with the phage's own replication — genes that benefit the host bacterium. These are called **moron genes** (from “more DNA” — extra genetic material that the phage has picked up from previous hosts) or, more formally, **auxiliary metabolic genes**. They can encode toxins that help the bacterium compete with other species, enzymes that allow it to metabolise new substrates, or proteins that protect it from environmental stresses. Some prophages carry genes that increase the bacterium's resistance to antibiotics, or that modify its surface structures in ways that evade the host's immune system. The prophage, in other words, can be an asset — providing the bacterium with capabilities it would not otherwise have.

This mutual benefit is the foundation of the **piggyback-the-winner** model, proposed by Ben Knowles and colleagues in a landmark 2016 paper in *Nature* [REF:knowles2016]. The model offers

an alternative to the kill-the-winner framework for high-density microbial environments like the gut. Knowles argued that, in ecosystems where bacterial density is high and encounters between phages and bacteria are frequent, the evolutionary advantage shifts from lytic predation to lysogeny. A temperate phage that integrates into a successful, abundant bacterial host “piggybacks” on the winner’s success — replicating for free every time the bacterium divides, spreading through the population without the costs of lytic replication. The host bacterium, in return, gains superinfection immunity and potentially useful accessory genes. The result is a mutualism: more microbes, fewer free viral particles, and a shift from predator-prey dynamics to a more cooperative relationship.

Knowles’ team tested this model by analysing 24 coral reef viromes and found that the ratio of virus-like particles to microbial cells decreased as microbial density increased — the opposite of what kill-the-winner would predict, but exactly what piggyback-the-winner predicts. In the dense microbial communities of coral reefs (and, by extension, the human gut), lysogeny appeared to be the dominant phage strategy.

The truth, in the gut, is probably that both models operate simultaneously. Kill-the-winner governs the dynamics of lytic phages, maintaining diversity by suppressing dominant species. Piggyback-the-winner governs the dynamics of temperate phages, which form stable associations with their hosts and spread through the population as silent passengers. The gut phageome is not a single-strategy system. It is a mosaic of lytic and lysogenic interactions, shifting between modes depending on bacterial density, environmental stress, and the specific phage-host pairings involved.

16.5 Phage-Mediated Horizontal Gene Transfer — Shuffling the Deck

Genes on the Move

In Chapter 2, we explored the idea that evolution is not only a story of vertical inheritance — parent to offspring, generation to generation — but also of horizontal gene transfer (HGT), the movement of genetic material between organisms that are not parent and offspring. We discussed how bacteria in the human genome had, over evolutionary time, contributed genes to our own DNA. Now we return to horizontal gene transfer, but at a different scale and with a different agent: phages as the vehicles.

Phage-mediated gene transfer — technically called **transduction** — is one of the three major mechanisms by which bacteria exchange genetic material (the other two being conjugation, which involves direct cell-to-cell contact, and transformation, which involves the uptake of free DNA from the environment). Transduction occurs as an accident of the phage replication process, and it comes in two forms.

Generalised transduction happens during the lytic cycle. When a phage replicates inside a bacterium, the host chromosome is typically degraded into fragments to provide raw materials for new phage genomes. Occasionally, the phage’s packaging machinery makes a mistake: instead of packing a copy of the phage genome into a new capsid, it packages a random fragment of the host bacterium’s DNA. The resulting particle looks like a normal phage — it has the correct protein coat and can attach to and inject DNA into a new bacterial cell — but the DNA it delivers is bacterial, not viral. The recipient bacterium receives a chunk of its neighbour’s genome and, if it can incorporate that DNA into its own chromosome, acquires whatever genes the fragment carried. The phage has inadvertently acted as a gene shuttle, ferrying bacterial DNA from one cell to another.

Specialised transduction is more targeted. It occurs when a prophage excises imprecisely from the bacterial chromosome during induction. If the excision cuts slightly outside the prophage boundaries, the departing phage genome takes with it a piece of the adjacent bacterial DNA. When this phage infects a new host, it delivers not only its own genes but also the bacterial genes it accidentally carried along. Because prophages tend to integrate at specific sites in the bacterial chromosome, specialised transduction tends to transfer the same neighbouring genes repeatedly — making it a powerful mechanism for spreading particular genes through a bacterial population.

A third mechanism, described more recently, is **lateral transduction**, which has been shown to be the most potent mode of phage-mediated DNA transfer yet discovered, capable of moving several hundred kilobases of bacterial genome in a single event — vastly more than the classical models predicted.

The Gene Shuttle in the Gut

The implications of transduction for the gut microbiome are profound. The human gut is, as we have established, one of the most densely populated microbial environments on Earth. It contains trillions of bacteria and a comparable number of phages, in close proximity, with phage infection and prophage induction occurring continuously. Every lytic cycle, every prophage excision, is an opportunity for transduction. The gut is, in effect, a vast horizontal gene transfer reactor — a molecular marketplace in which bacterial genes are being packaged, shipped, and delivered between species on a continuous basis.

What kinds of genes are being moved? Everything, potentially. But several categories stand out.

Antibiotic resistance genes are perhaps the most clinically concerning cargo. Transduction has been shown to move genes encoding resistance to beta-lactams, tetracyclines, vancomycin, and other antibiotics between bacterial species in the gut. The gut resistome — the collective pool of antibiotic resistance genes harboured by the gut microbial community — is maintained and expanded in part by phage-mediated transfer. This has obvious implications for the spread of antibiotic resistance, a topic we will revisit in detail in Volume 2 (Chapter 18). A particularly unsettling finding is that phage-plasmids — hybrid genetic elements that function as both phages and plasmids — have been identified carrying large numbers of antibiotic resistance genes, with one study cataloguing 60 phage-plasmids carrying a total of 184 resistance genes across diverse bacterial hosts.

Virulence factors — genes that enable bacteria to cause disease — are also frequent passengers. Some of the most notorious bacterial toxins are, in fact, encoded by prophage genomes rather than by the bacterium's own chromosome. The Shiga toxin produced by enterohemorrhagic *E. coli* O157:H7 — the toxin responsible for the severe kidney failure that makes this strain so dangerous — is carried on a prophage. Cholera toxin, the molecule that makes *Vibrio cholerae* lethal, is similarly encoded by a filamentous phage called CTXφ. Diphtheria toxin, botulinum toxin, scarlet fever toxin — the list goes on. In each case, the bacterium did not evolve the toxin gene on its own. A phage brought it. The boundary between “bacterial” virulence and “phage” virulence is, in molecular terms, blurred beyond recognition.

Metabolic genes — genes that allow bacteria to exploit new nutrient sources, resist environmental stresses, or produce useful compounds — are also transferred by phages. In the gut ecosystem, where competition for resources is intense, acquiring a new metabolic capability through phage-mediated transfer can provide an immediate competitive advantage. Phages, in this view, are not just killers and parasites. They are unwitting innovators — shuffling the genetic deck, dealing new

hands to the bacteria they infect, and accelerating the pace of microbial evolution in ways that no amount of spontaneous mutation could achieve.

CRISPR: The Bacterial Counter-Attack

We cannot discuss phage-mediated gene transfer without noting the bacterial immune system that evolved to defend against it — a system that has, somewhat unexpectedly, become one of the most important biotechnologies of the twenty-first century.

CRISPR — Clustered Regularly Interspaced Short Palindromic Repeats — is, at its core, a bacterial memory system for recording past phage infections. When a bacterium survives a phage attack, it can capture a small piece of the invading phage’s DNA and insert it as a “spacer” into a specific region of its own genome — the CRISPR array. This spacer serves as a molecular mugshot. If the cell encounters the same phage again, the spacer is transcribed into RNA, which guides a nuclease enzyme (typically Cas9 or a related protein) to recognise and cut any matching DNA — destroying the phage genome before it can replicate.

CRISPR-Cas systems are present in roughly half of all sequenced bacteria and are abundant in gut microbes, where the constant pressure from phage infection provides strong selective pressure for maintaining effective defences. The CRISPR arrays of gut bacteria contain spacers that match known gut phages, providing a historical record of past infections — a molecular diary of the ongoing phage-bacteria arms race within the gut ecosystem. Dion and colleagues, in a 2021 study, showed that CRISPR spacer sequences could be used to predict the bacterial hosts of unknown phages, providing one of the few tools available for linking the dark matter of the virome to specific bacterial targets [REF:dion2021].

The arms race does not stop at CRISPR. Phages have evolved **anti-CRISPR proteins** — molecules that neutralise the Cas enzymes, disabling the bacterial defence. Bacteria respond with modified CRISPR systems. Phages counter with new anti-CRISPRs. The cycle is continuous, rapid, and ancient — an evolutionary arms race that has been running for billions of years and shows no signs of resolution. It is also, incidentally, the arms race that gave us the CRISPR gene-editing technology that is transforming medicine, agriculture, and biotechnology. The molecular scissors that researchers use to edit human genes were originally forged in the war between bacteria and phages.

The Missing Players

We began this chapter by noting that the story of the human microbiome, as told in the preceding fifteen chapters, was incomplete. It was a story about bacteria and fungi — the visible inhabitants of the ecosystem — but it left out the invisible force that shapes their populations, moves their genes, and drives their evolution.

Phages are that force. They are the most abundant biological entities in the human body, as numerous as the bacteria they infect. They regulate bacterial diversity through kill-the-winner dynamics, preventing any single species from dominating the gut. They infiltrate bacterial genomes as prophages, conferring immunity and fitness advantages on their hosts while maintaining the option to reactivate and destroy them. They shuttle genes between bacterial species — including antibiotic resistance genes, virulence factors, and metabolic capabilities — reshaping the genetic landscape of the microbiome on a timescale of hours to days. And they have driven the evolution of one of the

most sophisticated defence systems in biology: CRISPR, a system so elegant that we have repurposed it for our own genetic engineering.

But we have not yet addressed what may be the most surprising aspect of phage biology: the relationship between phages and the human host. The phages in your gut do not merely interact with bacteria. They interact with you. They cross the epithelial barrier. They are found in the bloodstream. They engage directly with the immune system. And they may form part of a non-host immune defence — a viral shield in the mucus layer that protects epithelial surfaces from bacterial invasion.

That is the story of Chapter 17.

Chapter 16 — Key References

- [REF:mushegian2020] Mushegian, A.A. (2020). Are there 10^{31} virus particles on Earth, or more, or fewer? *Journal of Bacteriology*, 202(9), e00052-20. DOI: 10.1128/JB.00052-20
- [REF:shkoporov2019] Shkoporov, A.N. & Hill, C. (2019). Bacteriophages of the Human Gut: The “Known Unknown” of the Microbiome. *Cell Host & Microbe*, 25(2), 195–209. DOI: 10.1016/j.chom.2019.01.017
- [REF:twort1915] Twort, F.W. (1915). An investigation on the nature of ultra-microscopic viruses. *The Lancet*, 186(4814), 1241–1243. DOI: 10.1016/S0140-6736(01)20383-3
- [REF:dherelle1917] d’Hérelle, F. (1917). Sur un microbe invisible antagoniste des bacilles dysentériques. *Comptes Rendus de l’Académie des Sciences*, 165, 373–375.
- [REF:aswani2021] Aswani, V.H. & Shukla, S.K. (2021). An Early History of Phage Therapy in the United States: Is it Time to Reconsider? *Clinical Medicine & Research*, 19(2), 82–89. DOI: 10.3121/cmr.2021.1605
- [REF:norrby2008] Norrby, E. (2008). Nobel Prizes and the emerging virus concept. *Archives of Virology*, 153(6), 1109–1123. DOI: 10.1007/s00705-008-0088-8
- [REF:dutilh2014] Dutilh, B.E., Cassman, N., McNair, K., Sanchez, S.E., Silva, G.G.Z., Boling, L., Barr, J.J., Speth, D.R., Seguritan, V., Aziz, R.K., Felts, B., Dinsdale, E.A., Mokili, J.L. & Edwards, R.A. (2014). A highly abundant bacteriophage discovered in the unknown sequences of human faecal metagenomes. *Nature Communications*, 5, 4498. DOI: 10.1038/ncomms5498
- [REF:schmidtke2025] Schmidtke, D.T., Hickey, A.S., Engel, K., Tromans-Coia, C., Kimmerly, G.R., Johns, J.L., Forrest, K.M., Dutilh, B.E. & Bhatt, A.S. (2025). The prototypic crAssphage is a linear phage-plasmid. *Cell Host & Microbe*, 33(8), 1347–1362. DOI: 10.1016/j.chom.2025.07.004
- [REF:gregory2020] Gregory, A.C., Zablocki, O., Zayed, A.A., Howell, A., Bolduc, B. & Sullivan, M.B. (2020). The Gut Virome Database Reveals Age-Dependent Patterns of Virome Diversity in the Human Gut. *Cell Host & Microbe*, 28(5), 724–740. DOI: 10.1016/j.chom.2020.08.003
- [REF:minot2011] Minot, S., Sinha, R., Chen, J., Li, H., Keilbaugh, S.A., Wu, G.D., Lewis, J.D. & Bushman, F.D. (2011). The human gut virome: inter-individual variation and dynamic response to diet. *Genome Research*, 21(10), 1616–1625. DOI: 10.1101/gr.122705.111

- [REF:pargin2023] Pargin, E., Roach, M.J., Skye, A., Finlayson, B.J., Shkoporov, A.N., Egan, G.F. & Eddy, S.R. (2023). The human gut virome: composition, colonization, interactions, and impacts on human health. *Frontiers in Microbiology*, 14, 963173. DOI: 10.3389/fmicb.2023.963173
- [REF:allen2011] Allen, H.K., Looft, T., Bayles, D.O., Humphrey, S., Levine, U.Y., Alt, D. & Stanton, T.B. (2011). Antibiotics in feed induce prophages in swine fecal microbiomes. *mBio*, 2(6), e00260-11. DOI: 10.1128/mBio.00260-11
- [REF:knowles2016] Knowles, B., Silveira, C.B., Bailey, B.A., Barott, K., Cantu, V.A., Cobián-Güemes, A.G., Coutinho, F.H., Dinsdale, E.A., Felts, B., Furby, K.A., George, E.E., Green, K.T., Gregoracci, G.B., Haas, A.F., Haggerty, J.M., Hester, E.R., Hisakawa, N., Kelly, L.W., Lim, Y.W., Little, M., Luque, A., McDole-Somera, T., McNair, K., de Oliveira, L.S., Quistad, S.D., Robinett, N.L., Sala, E., Salamon, P., Sanchez, S.E., Sandin, S., Silva, G.G.Z., Smith, J., Sullivan, C., Thompson, C., Vermeij, M.J.A., Youle, M., Young, C., Zgliczynski, B., Brainard, R., Edwards, R.A. & Rohwer, F. (2016). Lytic to temperate switching of viral communities. *Nature*, 531(7595), 466–470. DOI: 10.1038/nature17193
- [REF:dion2021] Dion, M.B., Plante, P.L., Zufferey, E., Shah, S.A., Corbeil, J. & Moineau, S. (2021). Streamlining CRISPR spacer-based bacterial host predictions to decipher the viral dark matter. *Nucleic Acids Research*, 49(6), 3127–3138. DOI: 10.1093/nar/gkab133

Additional references to be verified and expanded during manuscript preparation.

Chapter 17: Phages, Immunity, and the Tripartite Relationship

In the previous chapter, we met the bacteriophages — the most abundant biological entities in the human body and, arguably, on the planet. We learned that they infect and kill bacteria, that they regulate bacterial diversity through kill-the-winner dynamics, that they integrate into bacterial genomes as prophages and ride along as silent passengers, and that they shuffle genes between bacterial species with a promiscuity that has profound consequences for antibiotic resistance, virulence, and microbial evolution. The picture that emerged was of a gut ecosystem shaped not by two players — human host and bacterial microbiome — but by three: host, bacteria, and the viruses that infect them.

But we left a question dangling at the end of Chapter 16, a question that would have seemed absurd to most virologists even twenty years ago. Can bacteriophages — viruses that, by definition, infect only bacteria — interact directly with the human body?

The answer, it turns out, is yes. And the interactions are far more extensive, more varied, and more consequential than anyone anticipated.

Phages cross the epithelial barriers of the gut, entering the body in vast numbers every day. They are found circulating in the bloodstream of healthy people — not as a sign of disease, but as a normal feature of human physiology. They engage directly with the mammalian immune system, triggering responses that range from anti-inflammatory suppression to full-blown antiviral activation, depending on context. And they appear to form a living antimicrobial shield in the mucus layer that lines every mucosal surface in the body — a non-host immunity, provided not by the body's own immune system but by the viruses that have taken up residence in its mucus.

This chapter is about those interactions. It completes the picture we began building in Chapter 16 and, in doing so, closes Volume 1. When we began this book, the microbiome was a community of bacteria living on and in the body. By now, it has become something far more complex: an ecosystem of bacteria, fungi, and viruses, communicating with the immune system, the brain, and the endocrine system, influencing metabolism, behaviour, and hormonal regulation, and interacting with the body at every level — from the molecular to the physiological. The phages are the last piece of this puzzle. They are also, perhaps, the most surprising.

17.1 Phage Transcytosis Across Epithelial Barriers

The Unexpected Commute

Consider the epithelial lining of the gut. It is a single layer of cells — remarkably thin, given its importance — that forms the boundary between the microbial world of the intestinal lumen and the sterile interior of the body. We discussed this barrier in Chapter 5 when we described the mucus layer that protects it and the tight junctions that seal the gaps between adjacent cells. The barrier is, in principle, a wall. Bacteria cannot normally cross it without causing infection. Immune cells patrol its inner surface, ready to respond if anything breaches the perimeter.

And yet, phages get through. Routinely. In enormous numbers.

The first clear demonstration came in 2017, when Sophie Nguyen, Jeremy Barr, and colleagues at San Diego State University published a study in *mBio* that showed phages are capable of **transcytosis** — active transport through living cells — across epithelial barriers [REF:nguyen2017]. The concept of transcytosis itself was not new. It is a well-established mechanism by which certain molecules and particles cross cell layers: they are taken up into a membrane-bound vesicle on one side of the cell (the apical surface, facing the gut lumen), trafficked through the cell's interior, and released on the other side (the basolateral surface, facing the body's interior). Antibodies use this mechanism. So do some nutrients. But nobody had shown that bacteriophages — complex viral particles, typically 50 to 200 nanometres in size — could do the same thing.

Nguyen's team demonstrated that they could. Working with confluent monolayers of epithelial cells grown in the laboratory — cell layers from the gut, lung, liver, kidney, and brain — the team applied phages to the apical (outer) surface and measured how many appeared on the basolateral (inner) side over time. The results were consistent across all cell types: phages crossed the barrier. The transport was **directional** — there was a significant preference for apical-to-basolateral movement, meaning the system was biased toward moving phages *into* the body rather than out of it. And it was not passive. The phages were not simply squeezing through gaps between cells. Microscopy revealed that phages were inside the cells — within membrane-bound vesicles, trafficking through the Golgi apparatus, the same intracellular sorting station that processes proteins and lipids destined for different cellular compartments. The phages were, in effect, using the cell's own delivery infrastructure to hitch a ride through the epithelial wall.

The rate was modest in percentage terms — approximately 0.1 per cent of applied phages were transcytosed over a two-hour period. But apply that rate to the estimated 10^{12} to 10^{13} phage particles in the human gut, and the numbers become staggering. Nguyen's team calculated that approximately **31 billion phage particles** cross the gut epithelial barrier into the body every day [REF:nguyen2017].

Thirty-one billion. Per day.

The Intrabody Phageome

This calculation demands a moment of recalibration. For most of this book, we have treated the body's interior — the blood, the organs, the tissues beneath the epithelial barriers — as essentially sterile. We have discussed the microbiome as something that lives *on* us — on mucosal surfaces, in the gut lumen, on the skin — not *in* us. The sterile-interior model is, of course, already under challenge: we discussed the controversial brain microbiome in Chapter 11. But the phage transcytosis data suggest something more sweeping. If 31 billion phages are entering the body every day through the gut alone — and similar transcytosis presumably occurs across the lung, the nasal passages, and other mucosal surfaces — then the body's interior is not sterile. It is suffused with bacteriophages. There is, in other words, an **intrabody phageome**: a population of phage particles circulating in the blood, accumulating in organs, and interacting with tissues throughout the body.

This is not a theoretical inference. Phages have been detected in the bloodstream of healthy individuals. A 2024 study by Lamy-Besnier and colleagues demonstrated that healthy human blood harbours a diverse viral community dominated by bacteriophages, with a significant fraction overlapping with gut virome sequences — consistent with the hypothesis that these circulating phages originated in the gut and translocated into the blood through epithelial transcytosis

[REF:lamybesnier2024]. The same study found that the blood phageome differed between healthy individuals and patients with Crohn's disease, suggesting that changes in the gut phageome — or changes in the gut barrier that alter phage translocation — could produce detectable signatures in the blood.

A companion study from the same year provided further mechanistic detail, proposing that phage translocation from the gut to the circulation involves two sequential barriers: first, epithelial transcytosis (the mechanism described by Nguyen), and second, filtration through the lymphatic system before phages reach the blood [REF:phageometransfer2024]. The lymphatic system, with its dense network of immune cells, likely intercepts a proportion of the translocating phages — and this filtration may itself be immunologically significant, as we shall see in section 17.2.

The discovery of the intrabody phageome raises questions that are, at this stage, more provocative than answerable. What are these phages doing inside the body? Are they merely inert particles being cleared by the immune system, or are they functionally active — infecting bacteria that have breached the gut barrier, modulating immune responses, or influencing tissue function? The honest answer is that we do not yet know. But the sheer scale of the translocation — billions of particles per day — makes it unlikely that the body has not evolved some form of response. And, as it turns out, it has.

17.2 Direct Interactions Between Phages and the Mammalian Immune System

A Virus by Any Other Name

The mammalian immune system evolved to detect and destroy viruses. It does not, in principle, care whether a virus infects human cells or bacterial cells — it recognises molecular patterns that are characteristic of viral particles in general: double-stranded RNA, unmethylated CpG DNA motifs, viral capsid proteins. These patterns are detected by the toll-like receptors (TLRs) and other pattern recognition receptors that we discussed in Chapter 12. When a phage enters the body — carried across the epithelial barrier by transcytosis, or introduced deliberately in the context of phage therapy — it presents these molecular signatures to the immune system. And the immune system responds.

The nature of that response is where things get interesting, because it is not a single response. It is a spectrum — ranging from aggressive inflammatory activation to quiet anti-inflammatory suppression — and the position on that spectrum depends on the specific phage, the dose, the route of entry, and the immunological context of the host.

The Innate Response: Recognition and Clearance

When phages enter the bloodstream — whether through natural transcytosis or deliberate injection — the innate immune system responds rapidly. Studies of intravenously administered phages in animal models have shown that more than 99 per cent of circulating phage particles are cleared from the blood within hours [REF:vanbellegheem2019]. The primary agents of this clearance are the tissue-resident macrophages of the liver (Kupffer cells) and the spleen, which engulf and destroy phage particles through phagocytosis — the same process these cells use to clear bacteria, dead cells, and other debris from the circulation.

The initial recognition involves several innate immune mechanisms operating in parallel. **Natural**

antibodies — pre-existing antibodies that circulate in the blood even before any specific exposure — can bind to phage capsid proteins and flag them for destruction. **Mannose-binding lectin (MBL)**, a soluble pattern recognition molecule, can bind to carbohydrate structures on phage surfaces. **Complement proteins** — the cascade system we described in Chapter 12 — can be activated by phage particles, coating them with opsonins that mark them for phagocytic uptake. Together, these mechanisms ensure that the vast majority of phages entering the bloodstream are rapidly intercepted.

But clearance is not the whole story. The process of recognising and destroying phages also generates immune signals — and those signals have downstream effects that extend well beyond the phages themselves.

The Anti-Inflammatory Side: Phages as Immunomodulators

In 2017, Jonas Van Belleghem and colleagues published a study that changed how immunologists think about phages [REF:vanbelleghem2017]. The team exposed human peripheral blood mononuclear cells (PBMCs) — the white blood cells circulating in the bloodstream — to purified preparations of *Staphylococcus aureus* and *Pseudomonas aeruginosa* phages. They meticulously controlled for bacterial endotoxin contamination (lipopolysaccharide, or LPS, the bacterial cell wall component that is a potent immune stimulant and a notorious contaminant of phage preparations), ensuring that the immune responses they observed were caused by the phages themselves, not by bacterial debris hitching a ride.

What they found was nuanced. The phages induced changes in the expression of hundreds of genes in the mononuclear cells, and the overall pattern was **predominantly anti-inflammatory**. There was strong upregulation of **IL-1RA** (interleukin-1 receptor antagonist) — a natural anti-inflammatory molecule that blocks the pro-inflammatory cytokine IL-1 — and marked suppression of **CXCL1** and **CXCL5**, chemokines that recruit neutrophils to sites of inflammation. The phages also induced expression of **IL-10**, a potent anti-inflammatory cytokine, and **SOCS3**, a suppressor of cytokine signalling that dampens inflammatory cascades.

At the same time, the phages did trigger some pro-inflammatory responses — modest increases in IL-1 β , IL-6, and TNF- α . But the balance was tilted firmly toward suppression rather than activation. The net effect was that phage exposure pushed the immune cells toward an anti-inflammatory state.

The mechanism appears to involve the **TLR3/TRIF signalling pathway**. TLR3 is a pattern recognition receptor that detects double-stranded RNA — a molecular signature typically associated with viral replication. When phages are taken up by immune cells and their nucleic acid is released inside the cell, phage-derived RNA (or RNA produced during intracellular phage gene expression) can activate TLR3 on the endosomal membrane. TLR3 signals through an adaptor protein called TRIF (rather than the MyD88 adaptor used by most other TLRs), and the TRIF pathway preferentially induces **type I interferons** (IFN- α and IFN- β) rather than the pro-inflammatory cytokines driven by MyD88 signalling. Type I interferons, in turn, suppress **NF- κ B** — the master transcription factor that drives inflammatory gene expression — creating a net anti-inflammatory effect.

Think of it as a molecular redirect. The immune system detects the phage, but instead of launching a full inflammatory assault (which would be appropriate for a human-tropic virus actively destroying host cells), it activates a pathway that produces interferons and suppresses inflammation. The response makes a certain evolutionary sense: phages are not a threat to human cells. They cannot infect them, cannot replicate in them, and cause no direct tissue damage. A full inflammatory re-

sponse would be wasteful and potentially harmful — collateral damage with no benefit. A measured interferon response, by contrast, puts the immune system on alert without inciting unnecessary inflammation.

The Pro-Inflammatory Side: When Phages Sound the Alarm

But the anti-inflammatory story is not the whole story. In 2019, Gogokhia and colleagues at the University of Utah published a landmark study in *Cell Host & Microbe* that revealed a very different face of phage-immune interactions [REF:gogokhia2019].

The team administered bacteriophages — targeting *Lactobacillus*, *Escherichia*, and *Bacteroides* — to germ-free mice (mice with no pre-existing microbiome). The phages, delivered in the absence of their bacterial hosts, had no bacteria to infect. They were, in effect, pure viral particles in a sterile gut. And they triggered a robust immune response. Specifically, the phage DNA activated **TLR9** — a pattern recognition receptor that detects unmethylated CpG motifs in DNA, which are common in bacterial and viral genomes but rare in mammalian DNA. TLR9 activation stimulated the production of **interferon-γ (IFN-γ)** — a potent pro-inflammatory cytokine that activates macrophages, promotes Th1 immune responses, and amplifies inflammation.

The clinical implications were stark. When Gogokhia’s team increased phage levels in mice with chemically induced colitis (a model for inflammatory bowel disease), the elevated phages **exacerbated the inflammation** — making the colitis worse, not better. The effect was dependent on TLR9 and IFN-γ: mice lacking either the receptor or the cytokine were protected from the phage-driven worsening. And when the researchers examined data from human patients, they found that patients with ulcerative colitis who responded successfully to faecal microbiota transplantation (FMT) had **lower** phage levels post-treatment than non-responders, and that mucosal IFN-γ levels correlated positively with phage abundance.

The Gogokhia study revealed that phages are not immunologically inert passengers. In the inflamed gut, they can act as amplifiers of inflammation — pouring fuel on an already burning fire. The mechanism is different from Van Belleghem’s anti-inflammatory pathway: Gogokhia’s phages activated TLR9 (which signals through MyD88, driving pro-inflammatory cytokines), while Van Belleghem’s phages activated TLR3 (which signals through TRIF, driving interferons and suppressing inflammation). The difference may reflect the form in which the phage is encountered — intact particles versus intracellular nucleic acid — or the immunological state of the host tissue. An inflamed gut, already primed for inflammatory responses, may respond to phage DNA very differently from a resting mononuclear cell in the blood.

The Tripartite Model

The picture that emerges from these studies — and from the broader literature on phage-immune interactions — is one of irreducible complexity. Phages do not have a single effect on the immune system. They have many effects, depending on context. In the blood and in resting immune cells, they tend to drive anti-inflammatory, interferon-dominated responses. In the inflamed gut, they can amplify inflammation through TLR9 and IFN-γ. In the mucus layer, as we are about to see, they provide a form of antimicrobial defence that does not involve the host immune system at all.

This complexity reflects the reality that the relationship in the gut is not binary — not simply “host and microbes” — but **tripartite**: host, bacteria, and phages, each interacting with both of the others. The host immune system responds to phages. Phages shape the bacterial community that the

immune system monitors. Bacteria carry prophages that can be induced by immune-mediated inflammation, releasing more phages that further modulate the immune response. The three-way interaction creates feedback loops that stabilise the ecosystem in health and can destabilise it in disease.

We noted a version of this theme in Chapter 16, when we discussed how phages shift from stabilising regulators to amplifiers of disruption depending on the ecological context. The immune data tell the same story from a different angle. The phage is the same. The molecule is the same. But its effect — anti-inflammatory or pro-inflammatory, protective or destructive — depends on the state of the system in which it operates. This is not a flaw in our understanding. It is a feature of the biology.

17.3 The BAM Model — A Non-Host Immunity Layer

The Viral Shield in the Mucus

We return now to the mucus layer — that gel-like barrier coating every mucosal surface in the body, from the gut to the lungs to the reproductive tract. We discussed it in Chapter 5 as a physical barrier separating the gut bacteria from the epithelial cells, a negotiated border maintained by goblet cells and constantly renewed. We noted that the inner mucus layer of the colon is essentially sterile — devoid of bacteria — while the outer layer is colonised by a specific community of mucus-adapted species. The mucus, in that account, was a passive structure: a gel that bacteria could or could not penetrate, depending on their capabilities.

In 2013, Jeremy Barr and colleagues proposed a radical reinterpretation. The mucus, they argued, is not merely a passive barrier. It is an actively defended zone — and the defenders are not immune cells but **bacteriophages** [REF:barr2013].

The proposal, published in the *Proceedings of the National Academy of Sciences*, was called the **Bacteriophage Adherence to Mucus** model, or **BAM**. Its central claim was simple but far-reaching: phages concentrate in mucus at higher densities than in the surrounding environment, and this enrichment provides an antimicrobial layer that protects the underlying epithelial cells from bacterial colonisation — a form of immunity that is not produced by the host but by the viruses living in its mucus.

The Evidence

Barr's team began by measuring phage-to-bacteria ratios on mucosal surfaces across a startling range of organisms — from cnidarians (corals and sea anemones, among the most ancient animals with mucosal surfaces) to humans. On every mucosal surface they sampled, the ratio of phages to bacteria was elevated compared with the surrounding environment. Mucus was enriched for phages. This was not a quirk of the human gut. It was a universal feature of mucosal biology, conserved across hundreds of millions of years of evolution.

The mechanism of enrichment involved a molecular interaction between phage capsid proteins and mucin glycoproteins — the large, heavily glycosylated proteins that form the structural backbone of mucus. Many phages carry **immunoglobulin-like (Ig-like) domains** on their capsid surfaces — protein folds that resemble the antigen-binding regions of mammalian antibodies. These Ig-like domains bind to the variable glycan (sugar) chains that decorate mucin molecules. The interaction

is weak at the individual level — a single Ig-like domain binding a single glycan chain is not a strong bond — but the multivalent nature of the interaction (many Ig-like domains on each phage capsid, many glycan chains on each mucin molecule) produces sufficient collective affinity to concentrate phages in the mucus layer.

The result is a living antimicrobial carpet. The phages embedded in the mucus are not dormant. They are functional — capable of infecting and destroying any bacterial cell that ventures into the mucus layer on its way toward the epithelium. In Barr's in vitro experiments, the presence of mucus-adherent phages reduced bacterial colonisation of epithelial cells and decreased epithelial cell death by bacterial pathogens. The phages were providing protection — not through any instruction from the host, not through any immune signal, but simply by being in the right place at the right time, doing what phages do: killing bacteria.

The Physics of the Hunt

Two years after the original BAM paper, Barr's group published a follow-up study in *PNAS* that explored the physical mechanism by which mucus-adherent phages encounter their bacterial prey [REF:barr2015]. The question was not trivial. Mucus is viscous — orders of magnitude more viscous than water. A phage particle embedded in mucus cannot swim (phages have no means of self-propulsion) and cannot diffuse freely (the mucin network impedes movement). How, then, does a phage in mucus find a bacterium to infect?

The answer was **subdiffusive motion**. Using high-resolution microscopy to track individual phage particles in mucus, the team found that mucus-adherent T4 phages (phages with intact Hoc proteins — the Ig-like domains that bind mucin) did not undergo normal Brownian diffusion. Instead, they exhibited subdiffusion — a slower, more constrained form of random motion in which the particle's displacement over time increases more slowly than in free diffusion. The phages were not stuck in place, but they were not freely floating either. They were tethered loosely to the mucin network, jiggling in place, able to explore a local volume but unable to travel far.

At first glance, this seems counterproductive. A phage that cannot move freely should encounter fewer bacteria, not more. But the mathematics of subdiffusion tell a different story. In a subdiffusive regime, particles spend more time in a given local area before wandering away. If bacteria are also present in that area — moving through the mucus on their way toward the epithelium — the prolonged local residence of the phage increases the probability of a chance encounter. The phage does not need to chase the bacterium. It just needs to be there when the bacterium arrives.

The experimental confirmation was elegant. When the team compared mucus-adherent T4 phages with mutant T4 Δ hoc phages (lacking the Hoc protein and therefore unable to bind mucin), the mutant phages diffused normally in mucus — moving faster and farther — but were dramatically less effective at killing bacteria. In microfluidic experiments using mucus-producing human cells and *E. coli*, wild-type T4 phages reduced bacterial colonisation of the epithelium by more than **4,000-fold** compared with the non-adherent mutant [REF:almeida2019]. The phages that moved less killed more. Adhesion to mucus, by constraining phage motion, paradoxically improved their antimicrobial effectiveness.

BAM as an Ancient Immunity

The BAM model proposes that this phage-mucus interaction is not accidental. It is a co-evolved mutualism — a symbiotic relationship in which both parties benefit. The phages benefit because

mucus is where their bacterial prey congregate. Bacteria must transit through the mucus layer to reach the epithelial surface, and a phage embedded in mucus is ideally positioned to intercept them. The host benefits because the phage layer provides an additional line of defence against bacterial invasion — one that does not require energy expenditure, immune activation, or the collateral tissue damage that accompanies an inflammatory immune response.

The conservation of this system across phyla — from cnidarians to mammals — suggests that it is genuinely ancient. Mucosal surfaces are among the oldest interfaces between animals and their microbial environments. If phages have been accumulating in mucus for hundreds of millions of years, then the BAM layer predates the adaptive immune system (which evolved in jawed vertebrates roughly 450 million years ago) and may even predate significant components of the innate immune system. It is, in this view, one of the oldest forms of antimicrobial defence on the planet — a viral shield that was protecting mucosal surfaces long before antibodies, T cells, or toll-like receptors existed.

In 2022, Barr's group published further evidence that this relationship is subject to natural selection. Phages exposed to mucus-producing cells over multiple generations **evolved enhanced persistence** in the mucosal environment — they accumulated mutations that improved their ability to adhere to mucus and maintain themselves in the mucin layer [REF:barr2022]. The phages were adapting to their mucosal niche, just as bacteria adapt to theirs. The BAM layer is not a static feature of mucosal biology. It is a dynamic, evolving system — shaped by the same evolutionary forces that drive every other biological interaction in the microbiome.

Limitations and Open Questions

The BAM model is elegant and well-supported by in vitro evidence, but several important questions remain. Most of the experimental work has been done in laboratory cell culture systems, not in living animals. The in vivo relevance — whether the phage-mucus layer genuinely reduces bacterial colonisation and infection in a living human gut, with all its complexity — has not been directly demonstrated. The model also assumes that mucus-adherent phages are functional and infectious, which is likely true for many but may not be true for all — some may be degraded, inactivated, or bound too tightly to interact with passing bacteria.

There is also the question of specificity. The BAM layer provides defence only against bacterial species for which a matching phage is present in the mucus. It is not a broad-spectrum barrier like the mucus itself or the innate immune system. If a pathogen arrives for which no mucosal phage exists — a novel invader, for example, or a species against which the local phageome has no representative — the BAM layer offers no protection. This is both its strength (exquisite specificity, no collateral damage) and its weakness (no protection against unfamiliar threats).

Despite these caveats, the BAM model represents one of the most conceptually important developments in microbiome science in the past decade. It reframes the mucus layer from a passive gel into an actively defended zone, adds a third dimension to mucosal immunity (alongside the physical barrier and the host immune system), and provides a concrete mechanism by which phages contribute to host defence — not through any direct interaction with host cells, but simply by doing what they have always done: killing bacteria.

17.4 Phages as Anti-Inflammatory Agents? Emerging Evidence

Beyond Killing Bacteria

Everything we have discussed so far in this chapter — transcytosis, immune modulation, the BAM model — describes phages interacting with the host as a consequence of their primary activity: infecting bacteria. The phage is not trying to cross the epithelial barrier, engage the immune system, or protect the mucus layer. It is trying to find and infect its bacterial host, and these other interactions are side effects — important, consequential side effects, but side effects nonetheless.

In this final section, we explore a more provocative possibility: that phages may have direct anti-inflammatory effects on host tissues that go beyond their antibacterial activity. This is a younger and more speculative area of research, but the early evidence is intriguing enough to warrant a careful look.

Clinical Observations

The earliest hints came from clinical observations in the phage therapy literature. Andrzej Górski and colleagues at the Ludwik Hirsfeld Institute of Immunology and Experimental Therapy in Wrocław, Poland — one of the few centres in Europe with a continuous tradition of phage therapy — published a retrospective analysis of patients treated with phage preparations for various bacterial infections [REF:gorski2012]. They noted that patients receiving phage therapy often showed reductions in inflammatory markers — C-reactive protein, white blood cell counts, erythrocyte sedimentation rate — that were more rapid and more pronounced than would be expected from bacterial killing alone. The observations were uncontrolled and retrospective, and the authors were appropriately cautious in their interpretation. But the pattern was consistent enough to suggest that something beyond simple antibacterial activity was going on.

Górski's group subsequently demonstrated, in a series of laboratory studies, that phage preparations could directly suppress **NF-κB** activity in mammalian cells [REF:gorski2017]. NF-κB is the master transcription factor that drives inflammatory gene expression — the molecular switch that, when activated, turns on the production of pro-inflammatory cytokines, chemokines, and adhesion molecules. Suppression of NF-κB is the mechanism of action of many anti-inflammatory drugs, including corticosteroids. The finding that phage preparations could achieve a similar suppression was unexpected and, if confirmed, potentially significant.

The mechanism remains under investigation, but two pathways have been proposed. First, as we discussed in section 17.2, phage RNA can activate TLR3, which signals through TRIF to produce type I interferons. These interferons, in turn, suppress NF-κB — creating an anti-inflammatory effect. Second, phage particles may directly interfere with the binding of bacterial LPS to TLR4 — the receptor that mediates the inflammatory response to bacterial endotoxin. If phages compete with LPS for TLR4 binding, or if phage-induced signalling through TLR3/TRIF suppresses TLR4/MyD88 signalling, the net effect would be a reduction in LPS-driven inflammation. This would be particularly relevant in conditions like sepsis, IBD, and metabolic endotoxaemia, where chronic low-grade exposure to bacterial LPS drives persistent inflammation.

The T Cell Connection

Phage interactions with the adaptive immune system add another layer. Górski's group reported that purified phage preparations could inhibit **T-cell activation** — reducing both proliferation and

cytokine production in human T cells stimulated in vitro [REF:gorski2017]. T-cell activation is a central event in adaptive immunity and in autoimmune inflammation. If phages can suppress it, even partially, the implications for inflammatory and autoimmune diseases would be substantial.

However — and this is a critical caveat — the Gogokhia study described in section 17.2 showed the opposite effect in a different context: phage DNA activated TLR9, driving IFN- γ production and T-cell activation that worsened colitis. The contradiction is not necessarily a contradiction. It may reflect the difference between purified phage proteins (which may suppress inflammation) and phage DNA (which may activate it), or between the immune environment of the blood (where anti-inflammatory responses predominate) and the immune environment of the inflamed gut (where pro-inflammatory responses are amplified). Context, once again, is everything.

Phages in the Treatment of Inflammatory Disease?

The possibility that phages have intrinsic anti-inflammatory properties has attracted interest from researchers working on inflammatory bowel disease, where conventional anti-inflammatory therapies (corticosteroids, anti-TNF biologics, JAK inhibitors) are effective but carry significant side effects and fail in a substantial proportion of patients.

The logic is seductive: a phage preparation could, in theory, provide a dual benefit — killing pathogenic or dysbiosis-associated bacteria through direct lytic activity *and* suppressing the accompanying inflammation through immunomodulatory effects. This would be a fundamentally different kind of anti-inflammatory therapy — one that addresses both the microbial trigger and the inflammatory response simultaneously.

The evidence, at this stage, is entirely preclinical. No randomised controlled trials have tested phage preparations as anti-inflammatory agents in human IBD or other inflammatory conditions. The in vitro and animal model data are suggestive but not yet sufficient to support clinical application. And the Gogokhia study serves as a sobering reminder that phage-immune interactions are bidirectional — the same phage that suppresses inflammation in one context can amplify it in another. Any therapeutic application would need to navigate this complexity with care.

This is, in other words, a hypothesis worth watching — not a conclusion worth acting on. The evidence is at the “preliminary / hypothesis-generating” level that we have flagged throughout this book when the science is promising but unproven. What is clear, even at this early stage, is that the old view of phages as biologically inert particles — relevant only insofar as they kill bacteria — is definitively wrong. Phages are immunologically active. They modulate the host’s inflammatory state. And the direction of that modulation depends on a set of variables — phage type, dose, route of exposure, host immune status, tissue context — that we are only beginning to map.

The Complete Ecosystem

We have now come to the end of Part IV, and with it, to the end of Volume 1.

Let us take stock of where we are.

We began this book, sixteen chapters ago, with a simple premise: you are not alone. The human body is host to trillions of microorganisms — bacteria, fungi, and viruses — that live on and in us, that we have carried since birth, and that participate in our physiology in ways that were, until recently, unsuspected.

In Part I, we met the microbiome: its scale, its diversity, its evolutionary origins, the story of how it colonises us across a lifetime, and the tools that scientists use to study it. In Part II, we took a guided tour of the microbial communities that inhabit every major body site — the gut, the skin, the lungs, the urogenital tract, the mouth, the fungal kingdom, and the controversial frontier of the brain. In Part III, we listened to the conversations between the microbiome and the body — the immune negotiation that begins at birth and continues for life, the gut-brain axis that links microbial metabolism to mood and behaviour, and the endocrine connections through which the microbiome participates in hormonal regulation from bile acids to oestrogen to the circadian clock.

And in Part IV — these last two chapters — we met the missing players. The bacteriophages. The most abundant biological entities in the human body, present in numbers equal to or exceeding the bacteria they infect. Regulators of bacterial diversity through predation. Architects of bacterial evolution through gene transfer. Silent passengers embedded in bacterial genomes as prophages, carrying the potential for both cooperation and betrayal. Travellers who cross the epithelial barrier in their billions, entering the body to circulate in the blood and accumulate in organs. Immunomodulators who engage the mammalian immune system in ways that range from anti-inflammatory suppression to pro-inflammatory amplification. And antimicrobial guardians who concentrate in the mucus that coats every mucosal surface, forming a viral shield that protects the epithelium from bacterial invasion — an ancient, non-host immunity that may predate every other form of antimicrobial defence in the animal kingdom.

The human microbiome is not a community of bacteria with a smattering of fungi and a background haze of viruses. It is a fully integrated, three-kingdom ecosystem — bacteria, fungi, and bacteriophages — in continuous, dynamic interaction with each other and with the human host. The host is not merely a habitat. It is a participant: shaping the microbiome through its immune system, its diet, its hormones, and its circadian rhythms, and being shaped in return through metabolites, neurotransmitters, and the constant molecular dialogue that we have traced across these seventeen chapters.

In Volume 2, we will turn from description to consequence. We will ask what happens when this ecosystem is disrupted — by antibiotics, by diet, by modern life — and what we can do to repair it. We will confront the promise and the limitations of probiotics, the radical ecology of faecal microbiota transplantation, and the emerging frontier of precision microbiome engineering. We will examine how microbial disruption contributes to inflammatory disease, metabolic disease, cancer, and neurological illness. And we will close by looking ahead — at the diagnostic, therapeutic, and philosophical questions that microbiome science is forcing us to reconsider.

But those are stories for another volume. For now, the ecosystem is mapped. The players are introduced. The conversations have been overheard.

The rest is what we do with the knowledge.

Chapter 17 — Key References

- [REF:nguyen2017] Nguyen, S., Baker, K., Padman, B.S., Patwa, R., Dunstan, R.A., Weston, T.A., Schlosser, K., Bailey, B., Lithgow, T., Lazarou, M., Luque, A., Rohwer, F., Blumberg, R.S. & Barr, J.J. (2017). Bacteriophage Transcytosis Provides a Mechanism To Cross Epithelial Cell Layers. *mBio*, 8(6), e01874-17. DOI: 10.1128/mBio.01874-17
- [REF:barr2013] Barr, J.J., Auro, R., Furlan, M., Whiteson, K.L., Erb, M.L., Pogliano,

- J., Stotland, A., Wolkowicz, R., Cutting, A.S., Doran, K.S., Salamon, P., Youle, M. & Rohwer, F. (2013). Bacteriophage adhering to mucus provide a non-host-derived immunity. *Proceedings of the National Academy of Sciences*, 110(26), 10771–10776. DOI: 10.1073/pnas.1305923110
- [REF:barr2015] Barr, J.J., Auro, R., Sam-Soon, N., Kassegne, S., Peters, G., Bonber, N., Hatter, M., Molineux, S., Guardado, A., Vega, R., Bailey, B., Pogliano, J., Danthi, P. & Rohwer, F. (2015). Subdiffusive motion of bacteriophage in mucosal surfaces increases the frequency of bacterial encounters. *Proceedings of the National Academy of Sciences*, 112(44), 13675–13680. DOI: 10.1073/pnas.1508355112
 - [REF:almeida2019] Almeida, G.M.F., Laanto, E., Asber, R. & Sundberg, L.R. (2019). Bacteriophage Adherence to Mucus Mediates Preventive Protection against Pathogenic Bacteria. *mBio*, 10(6), e01984-19. DOI: 10.1128/mBio.01984-19
 - [REF:barr2022] Barr, J.J. et al. (2022). Bacteriophages evolve enhanced persistence to a mucosal surface. *Proceedings of the National Academy of Sciences*, 119(27), e2116197119. DOI: 10.1073/pnas.2116197119
 - [REF:vanbelleghe2017] Van Belleghem, J.D., Clement, F., Merabishvili, M., Lavigne, R. & Vaneechoutte, M. (2017). Pro- and anti-inflammatory responses of peripheral blood mononuclear cells induced by *Staphylococcus aureus* and *Pseudomonas aeruginosa* phages. *Scientific Reports*, 7, 8004. DOI: 10.1038/s41598-017-08336-9
 - [REF:vanbelleghe2019] Van Belleghem, J.D., Dąbrowska, K., Vaneechoutte, M., Barr, J.J. & Bollyky, P.L. (2019). Interactions between Bacteriophage, Bacteria, and the Mammalian Immune System. *Viruses*, 11(1), 10. DOI: 10.3390/v11010010
 - [REF:gogokhia2019] Gogokhia, L., Buber, K., Roundy, J., Padgett, M.J., Hapfelmeier, S., Hooper, L.V. & Round, J.L. (2019). Expansion of Bacteriophages Is Linked to Aggravated Intestinal Inflammation and Colitis. *Cell Host & Microbe*, 25(2), 285–299. DOI: 10.1016/j.chom.2019.01.008
 - [REF:gorski2012] Górski, A., Międzybrodzki, R., Borysowski, J., Dąbrowska, K., Wierzbicki, P., Ohams, M., Korczak-Kowalska, G., Olszowska-Zaremba, N., Łusiak-Szelachowska, M., Kłak, M., Jończyk, E., Kaniuga, E., Golas, A., Purchla, S., Weber-Dąbrowska, B., Letkiewicz, S., Fortuna, W., Szufnarowski, K., Pawełczyk, Z., Rogóż, P. & Kłosowska, D. (2012). Phage as a Modulator of Immune Responses: Practical Implications for Phage Therapy. *Advances in Virus Research*, 83, 41–71. DOI: 10.1016/B978-0-12-394438-2.00002-5
 - [REF:gorski2017] Górski, A., Dąbrowska, K., Międzybrodzki, R., Weber-Dąbrowska, B., Łusiak-Szelachowska, M., Jończyk-Matysiak, E. & Borysowski, J. (2017). Phages and immunomodulation. *Future Microbiology*, 12(10), 905–914. DOI: 10.2217/fmb-2017-0049
 - [REF:lamybesnier2024] Lamy-Besnier, Q., Theodorou, I., Billaud, M., Zhang, H., Brot, L., Culot, A., Abriat, G., Sokol, H., De Paepe, M., Petit, M.A. & De Sordi, L. (2024). The human blood harbors a phageome which differs in Crohn's disease. *bioRxiv* / preprint. DOI: 10.1101/2024.06.04.597176
 - [REF:phageometransfer2024] (2024). Phageome transfer from gut to circulation and its regulation by human immunity. *bioRxiv* / preprint. DOI: 10.1101/2024.06.05.597592

Additional references to be verified and expanded during manuscript preparation.